SciDAC Visualization and Analytics Center for Enabling Technologies
Semi-Annual Progress Report
October 2007 through April 2008

E. Wes Bethel* and Chris Johnson†
Principal Investigators

Charles Hansen, Claudio Silva, Steven Parker, Allen Sanderson, Martin Cole‡

Sean Ahern, George Ostrouchov, Dave Pugmire, Jeremy Meredith§

Valerio Pascucci, Hank Childs, Peer-Timo Bremer, Daniel Laney, Ajith Mascarenhas, Kathleen Bonnell¶

Ken Joy, Christoph Garth, Bernd Hamann‖

Cecilia Aragon, Gunther Weber, Prabhat**

May 1, 2008

*Lawrence Berkeley National Laboratory
†Scientific Computing Institute, University of Utah
‡Scientific Computing Institute, University of Utah
§Oak Ridge National Laboratory
¶Lawrence Livermore National Laboratory
‖University of California – Davis
**Lawrence Berkeley National Laboratory

Contents

1 Executive Summary

The SciDAC Visualization and Analytics Center for Enabling Technologies (VACET) focuses on leveraging scientific visualization and analytics software technology as an enabling technology for increasing scientific productivity and insight. Our mission is to foster scientific insight through creating and deploying effective data understanding technology that is truly responsive to the needs of our stakeholders in the scientific research community who are "awash in data." It is widely accepted that one of the bottlenecks in contemporary science is the need to gain insight from vast collections of complex data.

The vision for our Center is to respond directly to this challenge by adapting, extending, creating when necessary and deploying visualization and data understanding technologies for our science stakeholders. Organized as a Center for Enabling Technologies, we are well positioned to be responsive to the needs of a diverse set of scientific stakeholders in a coordinated fashion using a range of visualization, mathematics, statistics, computer and computational science and data management technologies.

We are pleased to report accomplishments during the period of October 2007 through May 2008, both in terms of impact for scientific stakeholders and in terms of providing leadership in the visualization and analysis community.

VACET has made a substantial impact on the SciDAC community:

- VACET's production-quality, parallel capable Adaptive Mesh Refinement (AMR) visual data analysis software infrastructure (VisIt) was adopted as the software of choice by the SciDAC Applied Partial Differential Equations Center (APDEC) and the SciDAC Community Astrophysics Consortium. In the case of APDEC, this result has freed up project resources formerly allocated to maintain in-house AMR visualization software (see Section 2.1.)

- VACET researchers have developed a robust, quantitative analysis capability, applied to results from one of the candidate petascale applications (S3D, J. Chen, SNL-CA). This new capability provides the ability to see, for the first time, the evolution of combustion extinction pockets in large, time-varying 2D/3D combustion datasets (see Section 3.4).

- VACET researchers have developed a fundamental new analytic capability that enables quantification of the complexity of the burning regions and their evolution in time. In particular, the our scientific stakeholders are able, for the first time ever, to distinguish quantitatively the degree of turbulence in lean combustion simulations (see Section 3.3).

- Working with SciDAC accelerator scientists and our colleagues from the SciDAC SDM Center, VACET researchers have developed a novel approach for performing interactive visual data mining and analysis of large, complex multivariate data produced by advanced accelerator simulations. This capability allows science researchers to take advantage of large, parallel computational platforms to quickly identify interesting data, then perform visual analysis over all timesteps thereby gaining a better understanding of complex scientific phenomena, namely the process of Wakefield acceleration. This new capability was observed to have excellent scalability, as run on the XT4 system at NERSC (see Section 3.1).

- VACET ported VisIt to run on the Cray XT4/CNL architecture. This effort will allow scientists to effectively perform production-quality, large, parallel visual data analysis tasks on DOE's large Cray XT4 platforms located at LCF/ORNL and NERSC/LBNL (see Section 4).

VACET continues to set the standard for visualization research productivity and outreach:

- In this period, we report twenty-five peer-reviewed journal publications, eight conference proceedings publications, three invited articles, one book chapter, three posters, and eighteen submitted journal articles (see Section 7.1).
- VACET researchers won the IEEE Visualization 2007 Best Paper Award for the third year in a row. Additionally, of the five papers nominated for the Best Paper Award, three were from VACET researchers (see Section 7.4). We also received a Best Paper Award in the 2007 IEEE Symposium on Raytracing, and won the Best Poster Award at the 2007 IEEE Symposium on Visual Analytics Science and Technology.
- We have delivered over twenty different invited presentations (Section 7.2.)
- VACET members have performed significant outreach to the research community, highlighted by Ken Joy's service as the chair of the 2007 IEEE Visualization conference. In addition, team members served on many program committees and advisory boards, reviewed many papers, and participated in many workshops. (See Section 7.3.)

VACET also has been laying the groundwork for future results, both in terms of visualization research and in terms of deploying technology to our stakeholders:

- We have constructed a new algorithm for reconstructing material interfaces for Eulerian hydrodynamics simulations. We believe this algorithm can make a significant impact for our APDEC stakeholders in terms of increased accuracy for data analysis. (See Section 2.2.)
- We have made significant progress on implementing an efficient, parallelized streamline algorithm, which will benefit our APDEC and FACETS stakeholders. (See Section 2.3.)
- VACET is preparing an early release of 3D visual data analysis software infrastructure that will be part of the Earth System Grid's (ESG) Climate Data Analysis Toolkit (CDAT). The intent is to provide the new capabilities needed to facilitate analysis of data for the 5th assessment report of the Intergovernmental Panel on Climate Change (IPCC) (see Section 3.5).
- We have developed a novel visual data analysis technique, called the Equivalence Class Function (ECF), that combines statistical and visual data analysis. ECFs provide the ability to quickly see behavior and relevance of scientific phenomena in large, complex scientific datasets not possible with traditional visualization techniques. This technique, which originated from VACET's portfolio of basic research projects, has been integrated into the production-quality, parallel capable visual data analysis software infrastructure that is available at DOE's open computing facilities for broad use by the science community (see Section 5.2).
- We created a new technique for performing index/query operations, as part of the Query-Driven Visualization processing pipeline, that execute on the GPU and on multiresolution, temporally varying AMR data. This work is the first-ever implementation of index/query on the GPU, and first-ever implementation of hardware-accelerated temporal AMR visual data analysis (see Section 5.5).

2 Production-Quality, Parallel Capable AMR Visualization Software

2.1 Adoption by SciDAC Science Applications

VACET is a leader in production quality, parallel adaptive mesh refinement (AMR) visualization and analysis software infrastructure. We have recently deployed this infrastructure to our scientific stakeholders: the SciDAC Applied Partial Differential Equations Center (APDEC)[1] and the SciDAC Community Astrophysics Consortium. This result has numerous direct benefits to those researchers. First, it allows them to "buy rather than build", thereby resulting in a direct cost savings of scientific staff: they no longer need to develop and maintain AMR visualization software. Second, the VACET technology allows them to effectively use parallel computing infrastructure to perform interactive visual data analysis to help answer scientific questions in domains like combustion and astrophysics. Third, since the VACET technologies are deployed at DOE's open computing facilities as well as on the scientists' desktop, this result is an example of successfully bridging the gap across research, development and production deployment activities within DOE's science programs.

The adoption of VisIt by SciDAC stakeholders only took place because of a significant investment by VACET. This investment consisted of approximately nine months worth of effort by participants in the center employing expertise varying from performance optimization to handling of AMR data to graphical user interface design, as well as expertise with the VisIt project itself.

VACET's leadership role in AMR visualization is recognized by the broader visualization community. Last year, VACET researchers published multiple invited papers on the subject of AMR visualization, and gave multiple invited talks at international meetings on the subject of AMR visualization.

While having two major efforts within SciDAC adopt VACET technology as a "community standard" is a significant accomplishment, there is still a vast amount of work remaining. Current efforts focus on providing fundamental algorithms needed by stakeholders to enable scientific data analysis along with infrastructure enhancements aimed at easing use and production deployment. A summary of active projects in this space include:

- Embedded boundary analysis. See Section 2.2.
- Parallel streamline infrastructure. See Section 2.3.
- Multivariate, hardware-accelerated multivariate volume rendering. See Section 2.4.

2.2 Embedded Boundary Analysis

In many applications it is necessary to reconstruct or track the boundary surfaces, or "interfaces," between multiple materials that commonly result from multi-fluid Lagrangian-Eulerian hydrodynamics calculations. This problem, where the generated data sets have the characteristic that each cell contains a "volume fraction" representing the percentage of each material contained in the cell, now arises in a variety of applications, and is frequently called the embedded boundary problem.

[1] APDEC produces solver software, called the Chombo framework, and has in the past expended its own project resources to provide an AMR visualization application called ChomboVis. APDEC recently announced to its stakeholders that VisIt, the VACET AMR visualization software infrastructure, now supersedes ChomboVis as the tool of choice for visual data analysis of AMR datasets. See `http://seesar.lbl.gov/ANAG/chombo/chombovis.html` for more information.

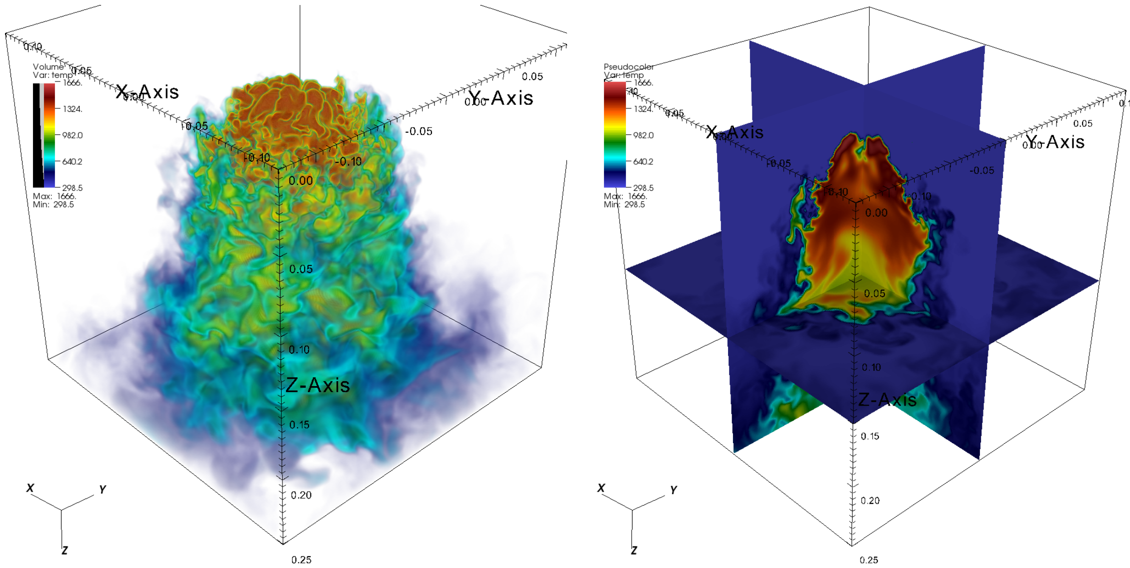

Figure 1: Production quality visualization of an AMR simulation of a hydrogen flame. Sample data courtesy of J. Bell and M. Day, Center for Computational Sciences and Engineering, LBNL.

The challenge is to utilize the material fractions in each cell to reconstruct the boundaries between materials. VACET researchers have been developing a new approach to this problem.

Before tackling the problem, we consulted with material interface reconstruction experts from Lawrence Berkeley, Lawrence Livermore, and Los Alamos. We strove to design a broad solution that could affect many stakeholders. Various laboratories require different solutions to this problem: some require a solver that takes only general volume fractions as input and outputs material boundaries. Another requires two-material solutions on AMR meshes. All require parallel implementations and the ability to work on diverse grids.

We have developed a new solver based upon active contours. Active contours are piecewise linear curves that move toward a specific target boundary. They work by balancing three forces, a force that attempts to keep the points of the curve equidistant, a force that attempts to limit the curvature of the curve, and a volume aware force that adjusts the curve motion according to volume constraints. Once an initial approximation of the curve is generated, the active contour quickly converges to the boundary. The model works on two- and three-dimensional data sets. Initial errors for these models are very good, several orders of magnitude lower than previous methods. The images below (Figure 2) show a comparison between the embedded boundary currently in VisIt and the new technique.

Material interface reconstruction algorithms are judged by many metrics: accuracy (in terms of preserving volume fractions), smoothness of output, quality of topology, size of output (in terms of number of primitives), and, of course, execution time. We designed our algorithm to perform well with these metrics. However, we plan to demonstrate its effectiveness before moving forward. If it is successful, we will extend the algorithm for widespread use by a broad community: adapting it for use on varying grid types and extending it for use on multiple materials. After that we will productize the algorithm and deploy it in VisIt.

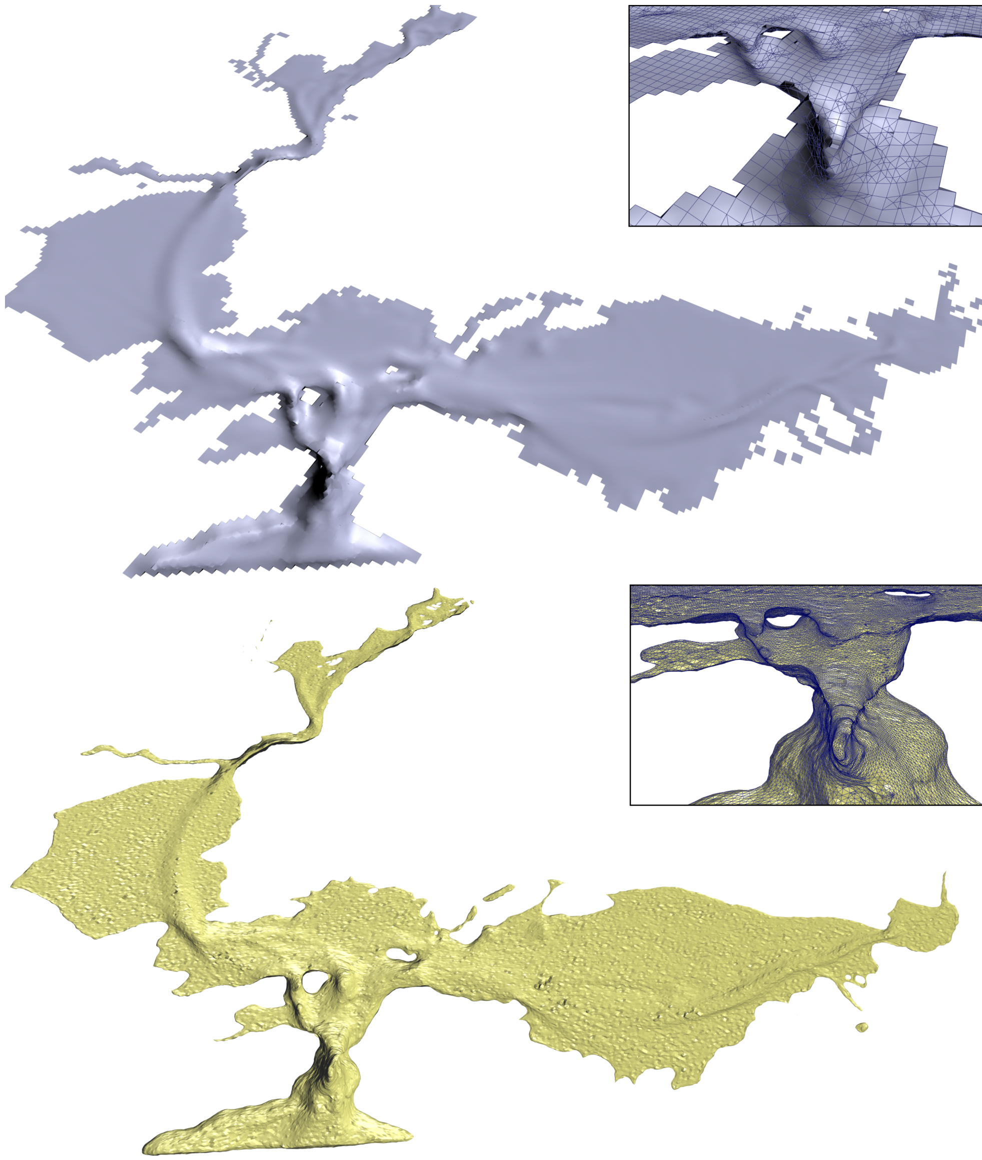

Figure 2: These images show an embedded boundary reconstruction of the San Francisco Bay Area (sample data courtesy of the SciDAC Applied Differential Equations Center). The top image shows a legacy algorithm implemented in VisIt. The bottom image shows the results of the new method, which exhibits a high degree of accuracy.

2.3 Parallel Streamline Infrastructure

There are several interrelated focus areas of this project. Broadly speaking, the high-level objective is to create software infrastructure for use in production-quality visual data analysis software tools that can effectively and correctly solve the vector field tracking problem in multigrid (AMR) contexts and that is suitable for use on parallel computing platforms. Several VACET team members from multiple institutions are collaborating on this effort, which is still in its early stages.

Our approach is to focus on developing a robust "streamline engine" that meets both use criteria – AMR/multigrid fields, and parallel implementation – then include this new engine into VisIt to achieve widespread use across a large base of scientific stakeholders. The AMR streamline capability

was initially requested by our APDEC stakeholders, though the robust, parallel implementation will be of significant value to stakeholders in many science areas (e.g., accelerator modeling, climate, combustion, fusion, etc.).

Preliminary work to date has focused on realizing the software infrastructure for a correct unigrid solution to be run on parallel platforms. Example output is shown in Figure 3, where the "streamline engine" is used to generate high quality "stream surfaces." We are presently benchmarking this implementation on various parallel platforms, and performing optimizations to improve efficiency. We expect a publication to result from this effort when completed. Near-term future work will focus on realizing the multigrid/AMR implementation, and we expect additional publications from that work. Meanwhile, we continue to interact closely with stakeholders in mathematics (APDEC) and fusion (J. Cary's FACETS project) to ensure our algorithmic development and implementation meet their scientific needs.

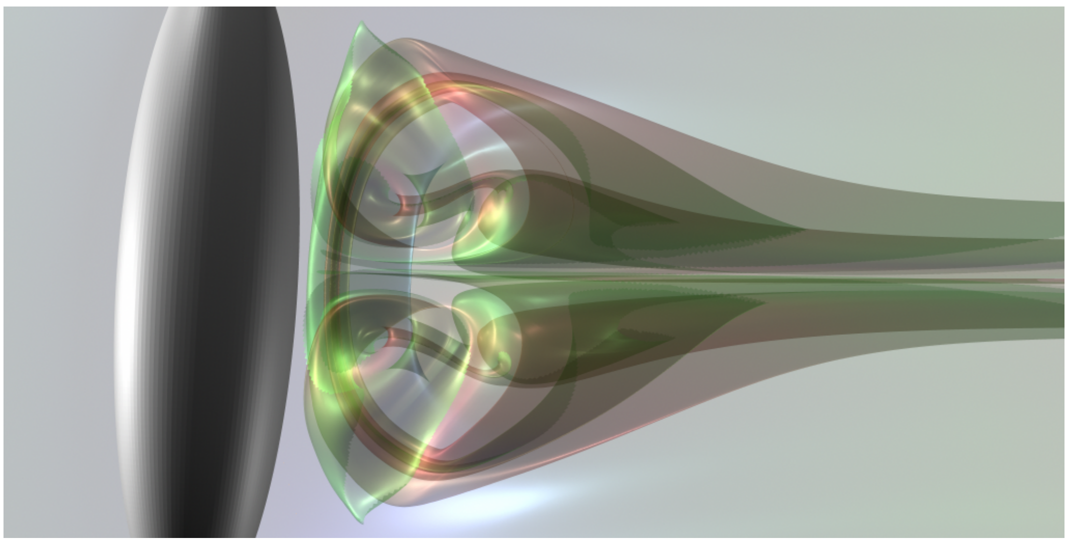

Figure 3: This image shows the high-quality stream surfaces generated by the new streamlines engine. Our long-term objective is to employ this crucial piece of fundamental technology into several different implementations of parallel vector field visualization techniques. The desired result is high-quality vector field visualization capability that is delivered, in production quality visualization software infrastructure, to a broad set of scientific stakeholders.

2.4 Multivariate, Hardware-Accelerated Volume Rendering

Although we have targeted VisIt as the production visualization tool for many of our partners, VisIt's volume rendering algorithm is below average in terms of functionality. We have decided to replace VisIt's algorithm with SLIVR: the SCIRun Library for Interactive Volume Rendering[2]. This project, which is a long term effort involving several VACET personnel from multiple organizations, has proceeded enough to produce an early prototype (see figure 4). SLIVR offers an immediate upgrade over the functionality available in VisIt (or any production visualization tool): it contains multi-dimensional transfer functions, which allow multiple scalar quantities to affect which parts of the volume are displayed. In addition, its lighting model is a significant upgrade over what VisIt previously deployed.

[2] See `http://slivr.sci.utah.edu`.

This effort will have long term benefit to nearly all VACET stakeholders. It is an excellent example of transitioning research prototype technology into production use for a wide audience.

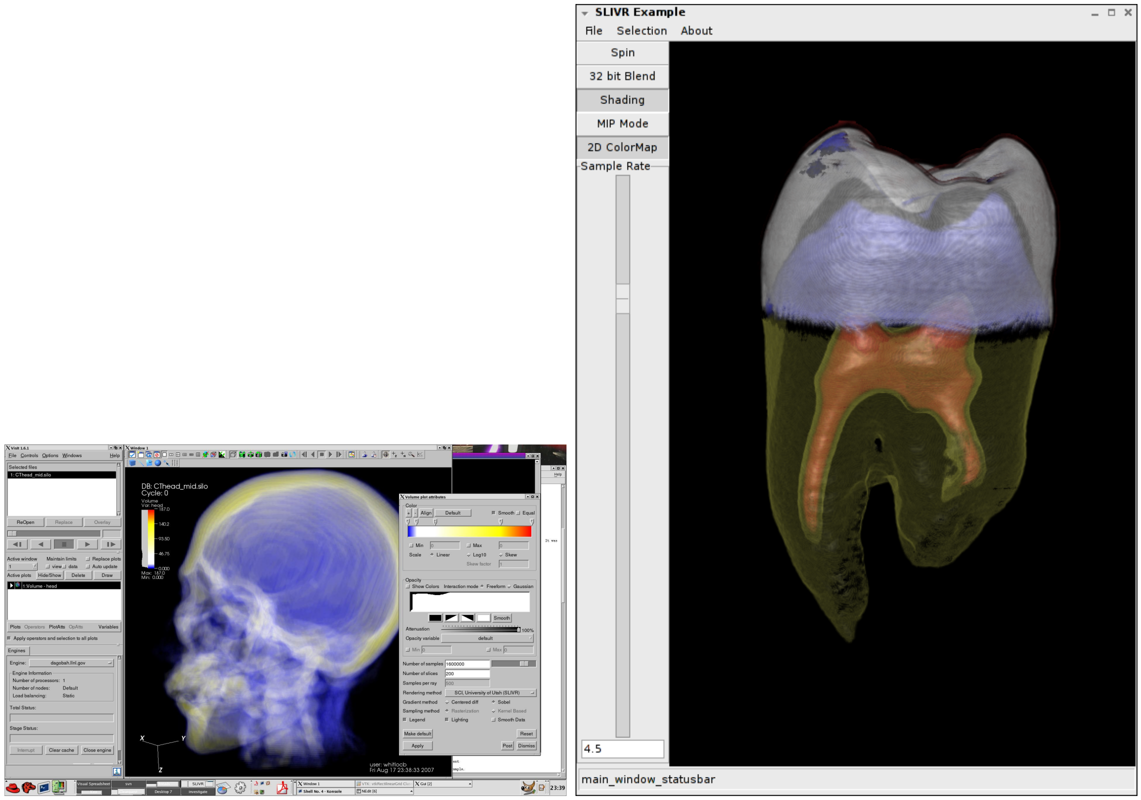

Figure 4: (Left image) Snapshot of image produced by an early SLIVR/VisIt integration. (Right image) Image of a tooth dataset rendered using a two-dimensional transfer function.

Future activities in this space (beyond productizing the deployment) focus on performance improvement, effective use on distributed memory platforms equipped with GPUs, streamlining the interfaces (user interface, software interface) between VisIt and SLIVR, and extending the fundamental SLIVR volume rendering algorithms for use on AMR grids.

3 Specific Stakeholder Projects

3.1 High Performance Visual Data Analysis for Advanced Accelerator Science

Working with advanced accelerator researchers from the SciDAC Community Petascale Project for Accelerator Science and Simulation (ComPASS) along with colleagues from the SciDAC Scientific Data Management Center, VACET researchers have developed a novel approach for rapid data mining of large, multivariate scientific datasets. This new capability responds to one of the central challenges in modern science, namely the the need to quickly derive knowledge and understanding from large, complex collections of data. Our new approach combines and extends techniques from high performance visual data analysis and scientific data management. This approach is applied to the problem of gaining insight from large, complex, time-varying datasets produced by a laser Wakefield accelerator simulation. Our approach leverages histogram-based parallel coordinates for both visual information display as well as a vehicle for guiding a data mining operation. Data

extraction and subsetting are implemented with state-of-the-art index/query technology. This approach, while applied here to accelerator science, is generally applicable to a broad set of science applications, and is implemented in a production-quality visual data analysis infrastructure. We conducted a detailed performance analysis and demonstrate good scalability on a distributed memory Cray XT4 system. The results of that performance study show: (1) we can perform interactive visual data mining on large (terabyte-sized) complex datasets on a distributed memory parallel platform; (2) this approach exhibits excellent scalability characteristics.

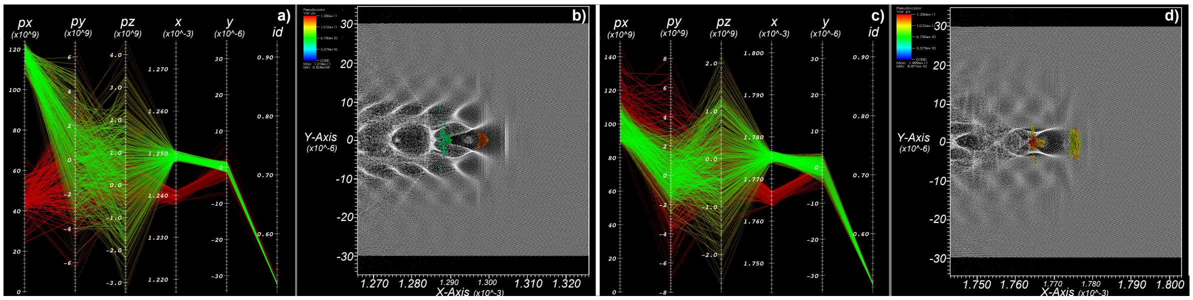

Figure 5: a) Parallel coordinates and b) pseudocolor plot of the beam at $t = 27$. Corresponding plots c,d) at $t = 37$. The context plot, shown in red, shows both beams selected by the user after applying a threshold of $px > 8.872 * 10^{10}$ at $t = 37$. The focus plot, shown in green, indicates the first beam that is following the laser pulse. In the pseudocolor plots b) and d), we show all particles in gray and the selected beams using spheres colored according to the particle's x-momentum, px. The focus beam is the rightmost bunch in these images. At timestep $t = 27$, the particles of the first beam (green in figure a) show much higher acceleration and a much lower energy spread (indicated via px) than the particles of the second beam. At later times, the lower momentum of the first beam indicates it has outrun the wave and moved into decelerating phase, e.g at timestep $t = 37$.

3.2 Astrophysics

VACET's strategy for providing help to the astrophysics community focuses on teaming with the SciDAC Science Application entitled Computational Astrophysics Consortium (CAC): Supernovae, Gamma Ray Bursts and Nucleosynthesis. In a nutshell, that project combines computational and experimental science to increase an understanding of the evolution of the universe by focusing on supernovae, gamma ray bursts and nucleosynthesis.

Our approach is to undertake efforts that provide visual and data analysis infrastructure that support both aspects of this large, ambitous astrophysics project. Our work with the simulation part of the community is outlined below in Section 3.2.1, and with the experimental community in Section 3.2.2. Both these communities are working together in the CAC: the experimental data provides validation of computational models, and the development of computational models helps guide how observations are performed.

3.2.1 Community Visualization Infrastructure

From a high level view, VACET's objective is to provide production-quality, parallel-capable visual data analysis software infrastructure for use by the SciDAC astrophysics community. To achieve that objective, we are working closely with "code owners" in the CAC to ensure that our primary production software, VisIt, meets their science-driven data understanding needs.

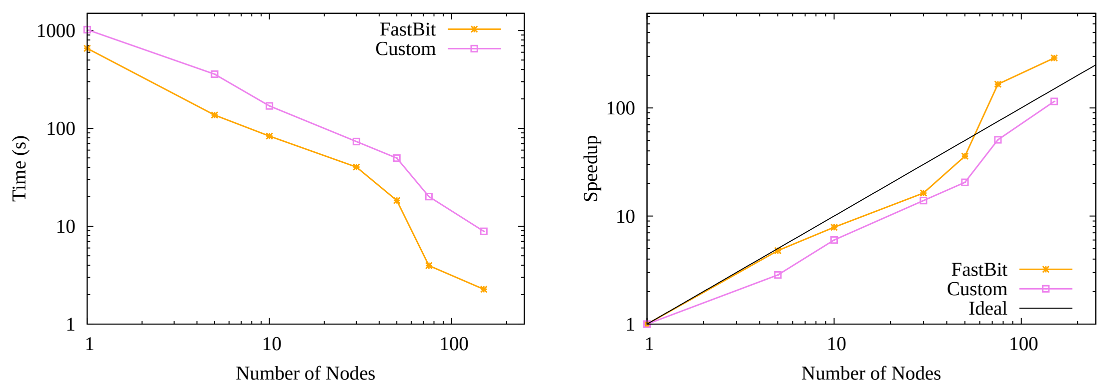

Figure 6: These images show the runtime and scalability speedup when performing queries on a large dataset (approximately 1TB in size) generated by a laser-wakefield accelerator simulation. We observe that queries on a 1TB dataset are answered in about two seconds when running 150-way parallel on a Cray XT4 (left image), and that the query response time scales well as we increase the level of parallelism (right image). The superlinear speedup evident in the scalability study (right image) is the result of being able to take advantage of local cache rather than having to go to the Lustre filesystem to obtain data.

The CAC is relying on John Bell, LBNL, to perform code development of a low mach number, incompressible flow model of a Type Ia supernova that includes radiation transport. Our existing work with APDEC in providing production-quality AMR visualization software (see Section 2) leaves us in very good position to quickly help this code team. At the April 2008 CAC All-Hands Meeting in Palo Alto, CA, John Bell told the CAC community that "VACET's VisIt is the tool of choice for performing visual data analysis of data from CASTRO" (CASTRO is his AMR-based code for supernova modeling). See Figure 7 for an example of this work.

VACET has also provided one-on-one support to other individuals in the CAC, namely Adam Burrows, in performing validation and exploration of data produced by the VULCAN code. Burrows is in the process of migrating from VULCAN to CASTRO for his work.

This project has a diverse and large set of needs that are documented in more detail on the VACET wiki at `http://www.sci.utah.edu/vacetwiki`. Near-term work plans focus on extending VisIt in a number of ways designed to meet a number of specific needs identified by the CCSE team at LBNL. As these needs are met, deployed in VisIt, and put into widespread use across CAC, we anticipate several things will happen. First, the needs list will grow longer and more focused on answering science-specific questions. These will be addressed in a way that achieves a good match between science priorities and VACET budget. Second, we expect our existing one-on-one relationships with CAC researcher sto mature and become even more fruitful. Through these relationships, we anticipate significant science results and discoveries to emerge.

3.2.2 Spectra and Light Curve Visual Data Analysis

Complementary to the simulation thrust within CAC is an experimental thrust. The CAC has relationships with ongoing experiments: the Nearby Supernova Factory (SNF), the Joint Dark Energy Mission (JDEM), the Supernova Acceleration Probe (SNAP), and the Large Synoptic Survey Telescope (LSST). The CAC uses an approach of validating simulation models against emperical observations, and these experiments all serve as sources of observations.

The visual data understanding needs of the CAC experimental group are documented on the

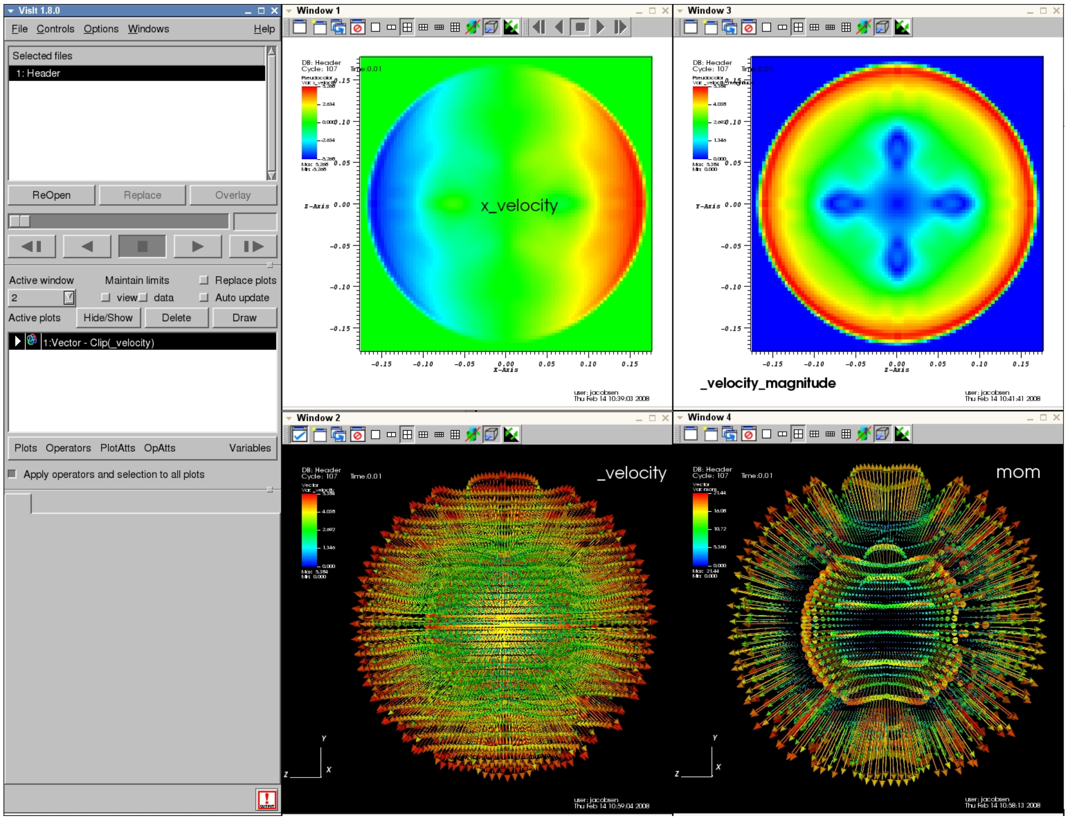

Figure 7: This image shows visual output of data generated by the CASTRO code, which is the AMR-based simulation in use by the CAC for modeling Type Ia Supernova explosions. The data was generated by Ann Almgren, LBNL, and the image created by Janet Jacobsen, LBNL.

VACET wiki. In brief, those needs can be summarized as follows: (1) replace "chi-by-eye" dataset comparison techniques with those that are more rigorous, insightful, and scalable; (2) determine a good set of parameters for capturing the population of spectral and light curves obtained through observation for the purpose of classification, matching, comparison and so forth; (3) better tools for performing temporal analysis and evolution of spectra, light curves and their associated parameters.

For (1), VACET has several interrelated efforts in this space, e.g., comparative analysis (Section 5.6) and equivalence class functions (Section 5.2). In addition, we have undertaken a separate project, called SpectraVis, that directly targets visual analysis and presentation of a large collection of spectra and light curves. An early prototype of that work, which was performed in collaboration with the Nearby Supernova Factory, provided a visual interface to browse a large collection of spectra and light curve data, and included some preliminary analysis. That work was well received by the science stakeholders. However, further development is currently in abeyance pending the arrival of a reference dataset from astrophysics collaborators in Europe.

While the SpectraVis work is on hold awaiting a validated collection of test data, we have focused on exploring new analysis and processing techniques that are useful for experimental astrophysicsts. We have developed a new technique for denoising noisy sky images, and experimented with applying fundamental analysis algorithms to particle accelerator data. We have also submitted research papers on the subject of effective data presentation and exploration methods useful to the experimental astrophysics community. In the long run, the intent is to simplify scientific knowledge discovery for the astrophysics community: this community is faced with trying to gain scientific

insight from data from diverse sources: collections of experimental/observed data and simulation results.

3.3 Structural Analysis of Time-Varying AMR Combustion Simulations

Working with John Bell and his team, VACET researchers have developed a novel feature detection, tracking and analysis capability. This capability allows, for the first time, combustion researchers to quantitatively study the relationships between turbulence, fuel mixture (lean vs. rich mixtures) and combustion efficiency. The approach is to use topological methods to identify "burning" and "non-burning" regions, and to perform a detailed structural analysis of those regions over time. The analysis includes studying the relationship between region characteristics (size, count, etc.) with other simulation variables (temperature, turbulence, etc.).

While initially applicable to Bell's 2008 INCITE simulation runs at NERSC (see Figure 8), the new methodology will be more broadly applicable to other science domains. Bell is PI on a SciDAC Science Application Partnership project (SAP) where his team will provide the primary code to be used by the SciDAC Community Astrophysics Consortium Science Application. The new VACET analysis capabilities will be crucial for enabling scientific knowledge discovery in large, complex AMR datasets emerging now from combustion and later from supernova modeling.

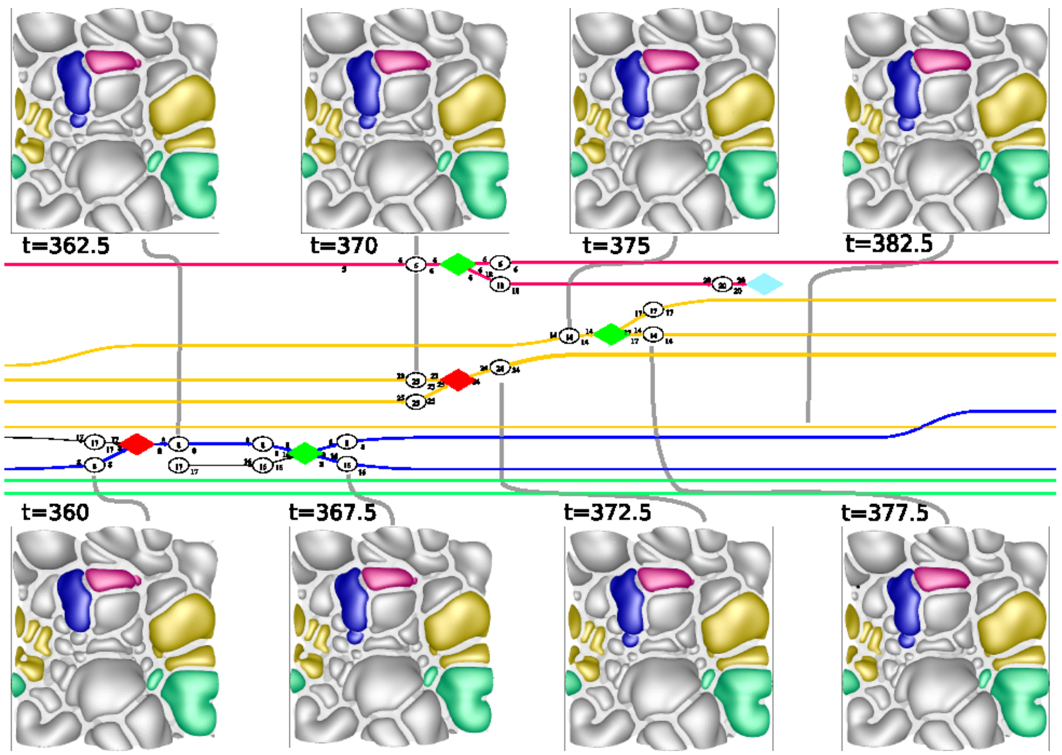

Figure 8: This image shows the results of the topological analysis of the AMR combustion simulation runs. Burning regions are identified using a topological basis, then their merging and splitting are tracked over time. The top and bottom row of images show identified regions, each of the regions is labeled with a unique color. The middle chart shows the temporal evolution of the splitting/merging of regions. This analysis was performed for several different simulation runs where various parameters like level of turbulence were varied from run to run.

3.4 Quantitative Analysis of Petascale Combustion Simulation Data

VACET researchers have developed new visual data analysis capabilities that help to accelerate scientific knowledge discovery in the petascale regime. In support of one of the candidate petascale applications (S3D, J. Chen, SNL), we have developed techniques that perform topological characterization of features in a combustion process, e.g., finding, tracking and analyzing extinction pockets, with an eye towards improving understanding the combustion process. Specifically, this analysis provides the ability to detect the dominant re-ignition mechanism associated with extinction pockets. This technique provides a robust tracking scheme and is of practical use for reliable quantitative analysis in large, time-varying 2D and 3D combustion simulation datasets (see Figures 9 and 10).

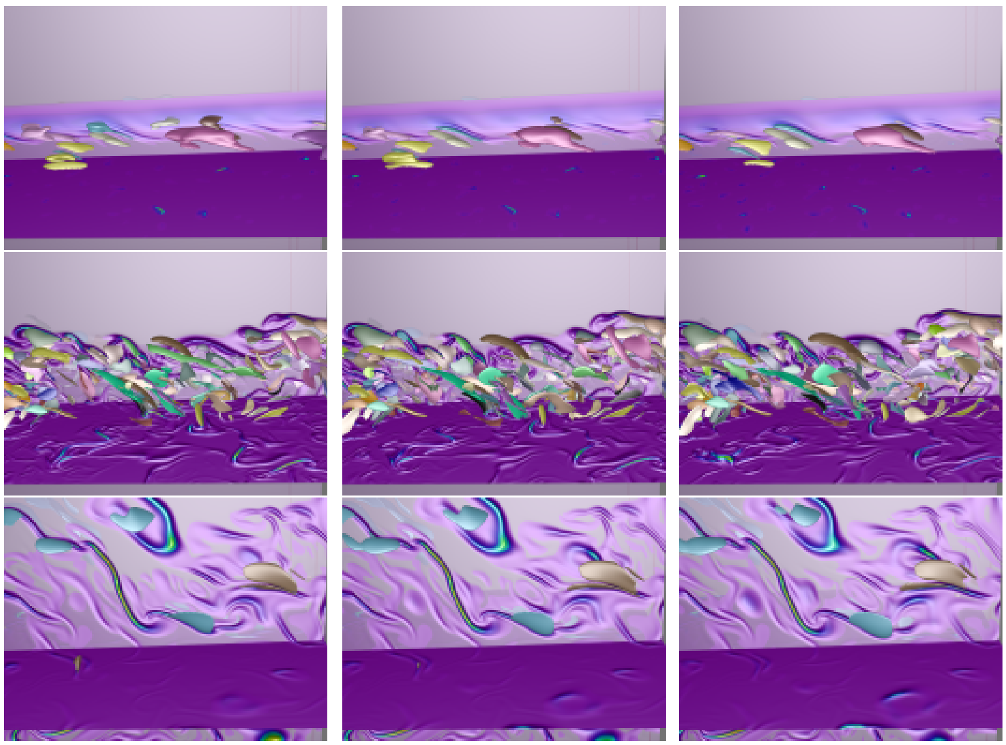

Figure 9: Frames showing tracked segments at three different temporal regions. In each frame, we place a vertical and horizontal 2D slice to provide context. We choose about 300 segments at $t = 75$ and track them forwards and backwards in time. Segments that split retain the color of their parent; segments that merge take on the color of one of their parents. Top image: segments at $t = 12, 14, 16$; middle: segments at $t = 72, 74, 76$; bottom: segments at $t = 190, 192, 194$.

3.5 Climate

The most advanced climate modeling systems seek to enable a new deeper understanding of the dynamics of global carbon cycle, atmospheric chemistry, land and ocean ecological processes and their coupling with climate. These new, advanced climate models will enable pursuing reliable answers to fundamental questions related to climate variability and global change at time scales ranging from decades to centuries. In this effort, VACET will work in close collaboration with the Earth Science Grid and provide new advanced data analysis and visualization tools to the CCSM Consortium and the climate modeling community in general. One target will be the deployment

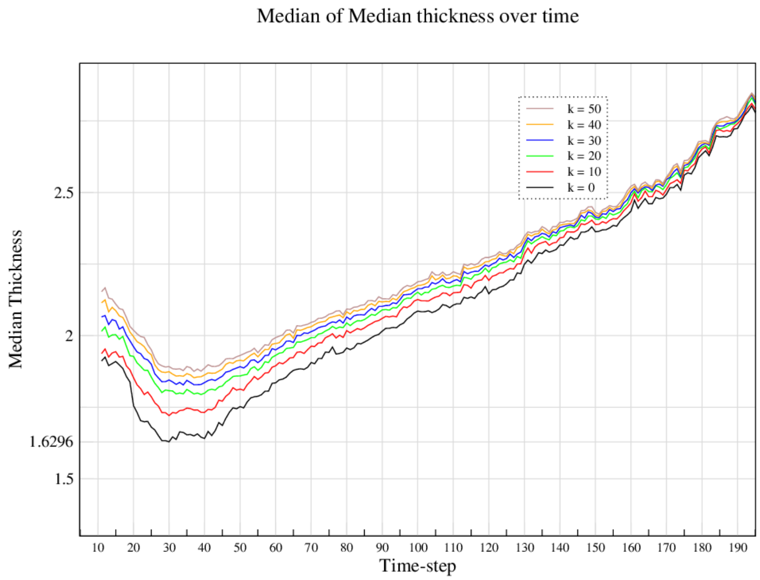

Figure 10: Quantitative analysis showing median thickness of combustion extinction pockets over time. A segment must have at least k vertices to be included in a count. This type analysis has never before been performed, and it shows an initial decline, then steady increase in extinction pocket median thickness over time.

of a first set of tools by the end of FY08, in time to facilitate the analysis of data for the 5th assessment report of the Intergovernmental Panel on Climate Change (IPCC)[3]. VACET is committed to supporting the needs of the Community Climate System Model (CCSM) Consortium[4] in collaboration with the Earth System Grid[5]. Our climate stakeholders include:

- Community Climate System Model Consortium, Primary contacts are John Drake (ORNL) and Phil Jones (LANL). Drake and Jones are co-PIs of the SciDAC Science Application entitled "Modeling the Earth System."
- Earth System Grid. Primary contact is Dean Williams, PI of the SciDAC Center for Enabling Technology entitled "Sharing a World of Data – Scaling the Earth Systems Grid to Petascale Data to enable faster, easier sharing of climate change research data."
- Primary target software tools: CDAT, VCDAT. New VACET capabilities will be layered atop the well established CDAT/VCDAT software infrastructure.

As this effort has a long time horizon, a summary of our objectives in this space are as follows: (1) deploy advanced visualization capabilities into the CDAT tool and create a clear path for similar integration in other tools; (2) extend the visualization software to incorporate domain specific requirements, data formats, and vector field visualization; (3) support time-dependent and cross-dataset comparison, visualization and analysis; (4) develop new analytic capabilities for

[3] http://www.ipcc.ch/
[4] http://www.cgd.ucar.edu/csm/
[5] http://www.earthsystemgrid.org/

climate data (first deployed into CDAT/VCDAT); (5) integrate with the VisTrails framework for tracking and logging of internal state and provenance information; (6) develop a visualization and data analysis scenario for understanding of complex coupled phenomena such as the multi-scale dynamics the complete carbon cycle on earth.

To date, the bulk of our work has focused on (1), which aims to extend the fundamental analysis and 2D visualization capabilities of the Climate Data Analysis Toolkit (CDAT), provided by ESG to its stakeholder community, with a baseline set of 3D visual data analysis tools. The basic idea is that by adding capability to CDAT, we will have a faster and broader positive impact on the climate community than if we ask them to migrate to a separate application (e.g., VisIt) for doing 3D/4D visual data analysis work.

In this work period, we have focused on software and release engineering issues that target fine-grained milestones associated with providing an alpha-quality release of software infrastructure that meets objective (1). In a nutshell, we are adapting and extending software infrastructure for use in the CDAT environment, as well as performing fundamental unit testing. Results of this work are shown in Figures 11 and 12.

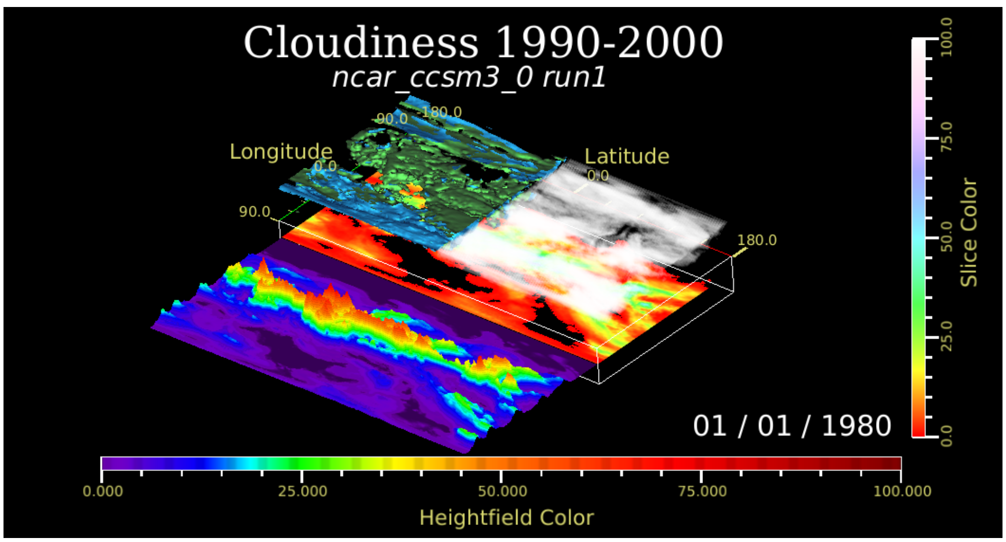

Figure 11: This image represents a new capability – a prototype integration of advanced visualization capabilities into the ESG-sponsored Climate Data Analysis Toolkit. This result proves the feasibility of the VACET approach to layering new capabilities onto high-quality software for use by a broad community.

We are exploring alternative scene layouts for data exploration and movie creation. For example, the image below (Figure 12) shows a possible layout including both embedded 3D data and slices on the side. This type of layout can increase both the realistic presentation of the information and the effective communication scientific the results. After completing our discussion with John Drake, Jamison Daniel, and the rest of the climate team at ORNL, we will improve the layout and content of the scene and generate a complete animation via scripting embeddable in CDAT. This capability will allow ESG to give the script to their users so they can produce similar presentations by simply changing the data source in the python script.

Next steps will take a more science-centric focus. The idea will be to team up with individual climate scientists and begin to apply the new 3D visual data analysis capabilities to specific science problems. The intent here is to better understand strengths and shortcomings of applying staple

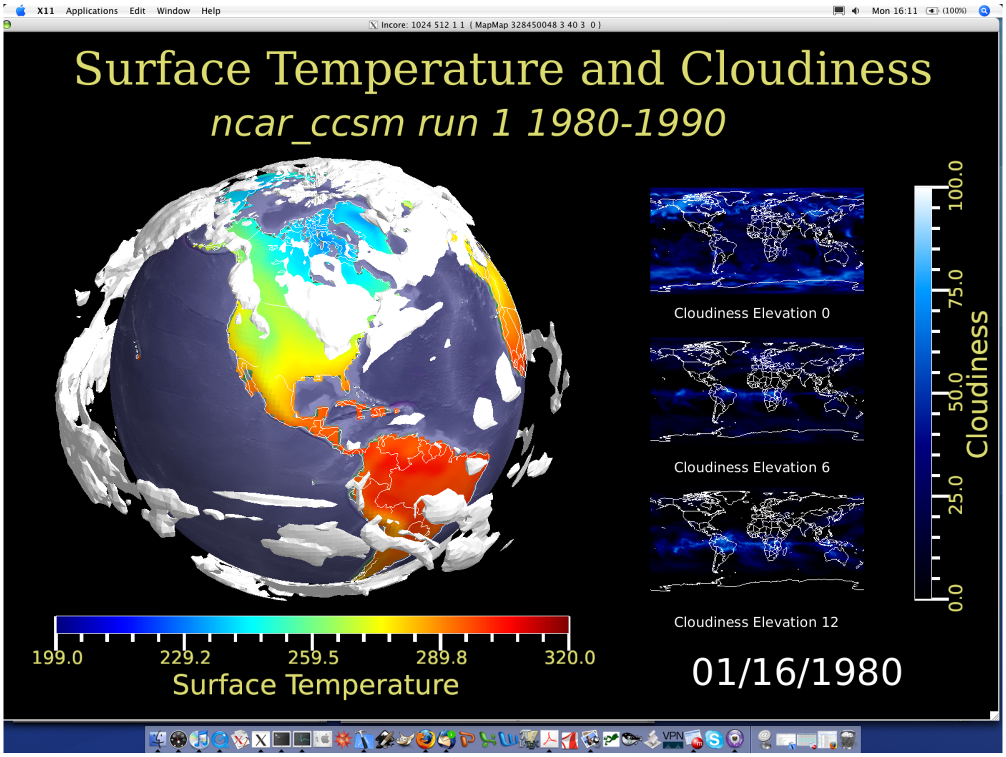

Figure 12: This image shows use of both traditional 2D and new 3D visualization capabilities within CDAT. This type of layout can increase both the realistic presentation of the information and the effective communication of the scientific results.

3D visual data analysis techniques (presented via CDAT) to specific scientific inquiries. We expect several different types of results from those interactions. First, we anticipate those interactions will provide priorities for next steps: will we better understand what additional 3D visual data analysis capabilities are needed by this community, and we will have a better understanding of scientific priority, which in turn will prioritize our work. Second, our work will have a direct, positive impact on ongoing climate science concerns, which is one of the primary VACET missions.

3.6 Fusion

3.6.1 Advanced Particle Visual Data Analysis

Many fusion projects use particle-based codes (e.g., GTC) to simulate scientific phenomena. As these codes are run on ever-larger machines, they "increase resolution" by using more and more particles to simulate scientific phenomena. As a result, contemporary runs can easily generate billions of particles per timestep. Data dumps of these temporally varying particle fields typically contain many variables at each particle – position, momentum, weight, etc.

Our fusion stakeholders have expressed the desire to be able to explore the nature of the particle orbits in an interactive manner. It is impractical to simultaneously view billions of particles due to the combination of extreme processing load and an insurmountable amount of visual clutter in the resulting images. Our approach, which was formulated in conjunction with our scientific stakeholders, is to focus visual analysis processing on "particles of interest," where "interesting" could be a compound boolean range query or a more elaborate set of statistical metrics. Displaying

other, non-particle fields, like potential or magnetic fields, in conjunction with particle-based visual data analysis presents a unique set of challenges. From a scientific standpoint, it would be useful to study the interaction of "interesting particles" with these fields. This combination leads to a number of investigation paths that our team is presently pursuing.

Previously, our effort has focused on visualizing the particle data in a time dependent fashion. As such, we developed the necessary infrastructure to enable image and movie generation. We and our stakeholders realized that static images showing a complete particle path over time was equally important. However, this objective was not immediately achievable due to the lack of internal and external software infrastructure to perform temporal slicing. We worked around the shortcoming by "building up" the time axis data as part of the movie production. This approach proved to be slow in execution and not particularly user-friendly. We tried a second approach, which holds more promise, of leveraging the latest version of the HDF5 I/O library (Version 1.8), which supports external links into readers and allows readers to scan across multiple datasets extracting the data across the time axis.

The second major accomplishment was the integration of tools for volume rendering the electric potential field data; during the simulations, particles move through and are influenced by this field. Previous implementations sliced the data using a brute force technique, where the underlying field uses unstructured prisms as its underlying mesh representation. Though unstructured, there is some underlying structure in how the prisms are assembled into a 3D mesh. After obtaining the scheme for how these codes write out an implicit semi-structured mesh of prisms, we were able to reconstruct the prism mesh and extend an existing volume rendering tool for use with this type of data. Additionally, we performed some internal optimizations of the prism mesh to result in faster rendering speed as compared to rendering the original prism mesh. Results from this work are shown in Figure 13.

These software prototypes provide the scientific stakeholders with two key items for post-query visualization. The first is to quickly view the path of a particle of interest. The second is to quickly view the affect of the potential field on the particle over time.

Next steps in this effort will focus on taking advantage of high performance index/query software infrastructure to accelerate the particle query process. Our team has performed some preliminary work in this area applied to Laser Wakefield Accelerator simulation data (see Section 3.1), and to explore leveraging traditional data analysis techniques in conjunction with visual data analysis. The overall objective is to accelerate the interactive visual data mining process for these large, complex and time-varying datasets.

3.6.2 Production Quality Fusion Visual Data Analysis Software Infrastructure

One of our fusion stakeholders, the SciDAC Fusion Science Application Framework Application for Core-Edge Transport Simulations (FACETS) project led by J. Cary, Tech-X Corporation, has identified a need for production quality visualization for fusion and accelerator simulation modeling capability. They wish to migrate from a collection of older commercial and unsupported tools, which collectively meet only part of their scientific data understanding needs, to a single modern toolset that meets their needs across multiple disciplines. They have chosen to work with VACET and plan to adopt VisIt as their community-wide visual data analysis infrastructure. This approach allows us to focus on collecting needs from community representatives, then deliver results that will have impact on a broad scientific community.

Specific accomplishments this period include:

- *Improved Streamlines.* This work, which is described in more detail in Section 2.3, combines

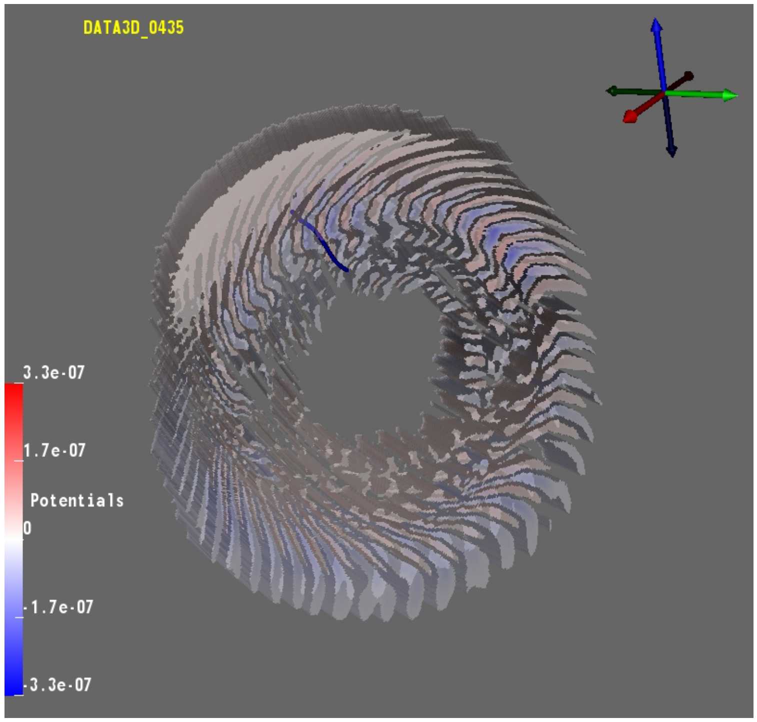

Figure 13: Image of a particle as it moves through the surrounding electrostatic potential. The particle's interaction with the potential results in a large radial diffusion. By using a transparent rendering, it is possible to "see through" the electrostatic potential. Also, potential values near zero are removed, which aids in viewing the structure of the mode waves within the field.

the effort of several VACET team members in support of multiple projects. In the case of the fusion stakeholders, the objective is provide the ability to compute and display streamlines (and advecting particles) that evolve over time and on temporally varying flow fields.

- *Poincaré plots.* The team successfully modularized a capability that formerly existed only within SCIrun and have demonstrated its use in a standalone fashion. The longer term target is to migrate this capability into VisIt for widespread use by a larger community of fusion stakeholders.

- *Parallel GTC Data Reader.* In response to needs from our GTC fusion stakeholders, we developed a robust, parallel GTC data reader. This reader has been deployed in VisIt, and is available to the world in versions 1.8 and later.

- *Data I/O Usability Improvements.* Our fusion stakeholders at Tech-X would like to specify a subset of database readers (file loaders) that are of interest. This capability greatly streamlines Tech-X's work flow when a particular file type suffix is shared by multiple file formats.

Future work in this space includes:

- *Poincaré plots.* Port the standalone tool into VisIt for widespread deployment.

- *Further Data I/O Usability Improvements.* Provide for much more flexible I/O options and operations: e.g., subset loading, dynamic (e.g., load time) assignment of variables in a file to different variable or coordinate fields, and so forth. This need reflects Tech-X's desire to

move towards using the Fusion Simulation Markup Language (FMSL) as an xml-interface to a "zoo" of different fusion data file formats.

- *Temporal visualization.* In particular, particle advection and streamline computation on unsteady flows.
- *Comparative Visualization.* Direct comparisons of data from multiple and different simulation runs, as well as comparisons between simulation and experimental data.

4 Common Infrastructure Projects

4.1 Port VisIt to Cray XT4/CNL

One of VACET's primary mission objectives is to provide production quality, parallel capable visual data analysis software infrastructure for use at DOE's open computing facilities. This past period, the VACET team ported VisIt to run in parallel on the Cray XT4 CNL architecture. This port was performed at LCF/ORNL, and tested at both LCF/ORNL and NERSC/LBNL. An image produced by one of these tests in a 256-way parallel run is shown in Figure 14.

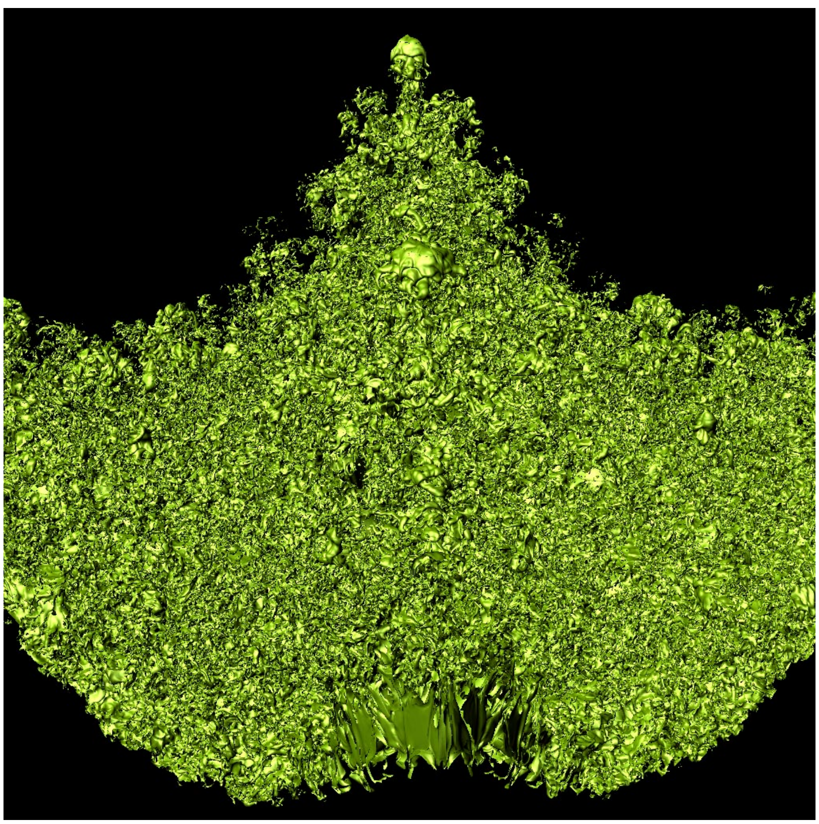

Figure 14: This image was generated by VisIt running on 256 nodes of Jaguar, a Cray XT4 system located at LCF/ORNL. The data, which was produced by a Richtmeyer-Meshkov simulation, consists of eight billion zones and 960 domains. The VisIt client ran on a Linux desktop machine at ORNL. Detailed scaling studies and additional information will be presented at the Cray Users Group meeting in May 2008.

4.2 Data-Parallel Infrastructure for Statistical and Analytical Computing

It is not possible to browse a petabyte of data. To see 10% of a petabyte of data at the rate of 100 MB per second takes 35 8-hour days! At this rate we are clearly exceeding the limits of of human bandwidth and patience. Although simulation brings with it the notion of unlimited transparency of simulated systems, as data sizes approach a petabyte, the difficulty of browsing such vast amounts of data reintroduces a great deal of opacity. There is a need for analysis methods that assist the human scientist in exploration of data by organizing the data and suggesting (or finding) interesting behavior. GTC particle simulations can easily produce petabyte data sets although present output is usually limited to a random subset of the simulated particles. A particle selection capability is needed at two stages: in the simulation code itself to limit the amount of I/O and further during visualization to provide interactive rendering speed of the data that has been output.

In addition, fusion scientists and climate scientists can not routinely do interactive statistical analyses on data sets much larger than a gigabyte simply because statistical software is largely serial. The methods needed include various analyses of variability, empirical orthogonal functions, and analyses of correlation or dependence together with graphical presentation of the resulting maps.

To answer the above application needs, we are developing a runtime environment for data-parallel statistical computing with R. The R system for statistical computing is the *de facto* standard for interactive statistical analysis and graphics and prototyping of new statistical algorithms. It is licensed under GPL and already provides opportunity for parallelism through its Rmpi and snow packages. Within this data-parallel statistical computing environment we will deploy the statistical analysis and graphics needs of Fusion and Climate Modeling.

Early work in this space focuses on the development and testing of infrastructure (file loaders, etc.) to be used in conjunction with a data-parallel implementation of statistical analysis tools based upon the R package. We are initially targeting stakeholders in Fusion (GTC), but will later expand scope to include stakeholders in climate. In the longer term, our objective is to package and release the new parallel execution infrastructure to be part of the R package (called *paRallel*).

5 Technology Incubation Projects

5.1 Uncertainty Visualization

An important visualization research problem is how to effectively convey uncertainty information along with traditional visual data representations. Our work investigates the problem from a graphical data analysis standpoint. By using descriptive statistics to summarize both characteristic features of a data distribution and measures of uncertainty, we can achieve a more cohesive understanding of the information. In this work, we reexamine the box plot and its relatives and develop a new hybrid summary plot that combines moment, cumulant, and density information along with higher order descriptors that rely on distribution fitting. In view of the important role summarizing plots has in decision making, our work focuses on using advanced visualization techniques to incorporate additional descriptive parameters, while simultaneously improving the comprehensibility of summary plots.

Uncertainty information is an important characteristic associated with much of the data scientists encounter. Such information is typically included with the data as 2D charts and graphs, however incorporating uncertainty into visualization techniques has proved quite challenging. We are examining a class of uncertainty data that is characterized as a set of probability density functions (PDFs) defined across a triangular mesh, and explores ways at visually presenting this

data.

During this work period, we have conducted initial investigations into these fundamental issues and have submitted two journal papers describing early results.

In the near future, our efforts will focus on two related areas. First, we are investigating direct high dimensional probabilistic data visualization with applications in fusion simulation and biophysical phenomena. Second, we are exploring general methods for uncertainty visualization in scalar and vector data fields.

5.2 Equivalence Class Functions

VACET Researchers have developed a new class of derived quantities – Equivalence Class Functions (ECFs) – designed to greatly expand the ability of end users to explore and visualize data. These functions are defined over equivalence classes (i.e., groupings) of elements from the original mesh, and produce summary values for the classes as output. ECFs can be used in the visualization process to directly analyze data, or can be used to synthesize new derived quantities on an original mesh. The design of ECFs enable a parallel implementation that allows the use of these techniques on massive data sets that require parallel processing. When integrated and used in concert with standard visualization techniques, ECFs offer a powerful way to manipulate, visualize and analyze scientific data. Formal definition of ECFs and their efficient data-parallel implementation for data sets at the extreme scale will extend the application of many statistical methods to data set sizes they previously could not reach. These concepts have been implemented into VisIt, our production-quality, parallel capable visual data analysis software infrastructure. As such, this technique is immediately usable by a broad range of scientific stakeholders on virtually all modern HPC platforms.

Future plans in this area will focus on applying the techniques to large and complex scientific datasets as part of our one-on-one interactions with science stakeholder teams. We expect those interactions will lead to the design and development of extensions to this fundamental technique.

5.3 Function Fields

As computational science evolves, simulations produce output data of increasing complexity. Traditionally, simulations (and experiments) generate scalar and vector fields over a computational domain. In relatively recent times, simulations have begun to undertake substantially more advanced modeling where they use a "function field" at each point in the computational domain rather than scalars or vectors.

For example, radiation transport problems collect data in "bins," representing the data as histograms at each data point. Most visualization methods are based upon the display of scalar fields, and these methods cannot be lifted to develop similar methods for function fields. Our objective is to develop, together with our science stakeholders, a number of visualization techniques that allow the exploration of function fields.

Function fields map points in $n-$dimensional Euclidean space to one-dimensional scalar functions. These functions arise in neutron transport applications, hyper-spectral imaging applications, and air-pollution control modeling, among others. In general, visualization of these methods have been based upon probe-based methods that are used to derive scalar fields to show the similarity structure of function fields. Multiple probe methods allow more information to be seen and provide users with greater visualization flexibility over time and space.

To date, work in this space has focused on fundamental research of the technique applied to hyperspectral image data, along with associated publications. Future work will target science stake-

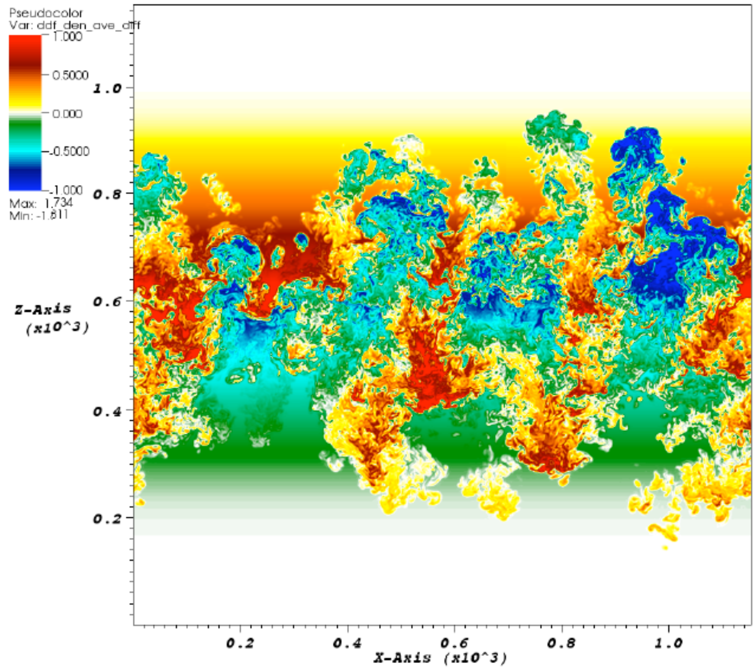

Figure 15: The absolute value of difference from ECF-synthesized average density (averaged over xy-planes). The large blue region in the image represents an unusually large, coherent volume of light fluid that has moved an unusual distance into the heavy fluid, relative to other regions in the dataset. This type of scientific insight is not possible with traditional visualization techniques.

holders, primarily in astrophysics, who are developing codes that model supernova explosions using a combination of AMR, CFD, MHD, reactive chemistry and radiation transport. As the techniques for function field visual data analysis and data management mature, they will be integrated into VisIt for widespread deployment.

5.4 Parallel, Out-of-Core Visualization Using Hybrid Resources

This project, which is part of VACET's basic research portfolio, focuses on leveraging multiple GPUs, to create visualizations of very large collections of data. We are working with John Owens, UCD, who is part of the SciDAC Institute for Ultrascale Visualization on this project. Our aim is to produce the infrastructure to render very complex, out-of-core scenes using unbiased global illumination. The system facilitates the rendering of hundreds of millions of triangles, gigabytes of texture, and complex shading by forcing high levels of coherency. A hybrid processing paradigm is created, using both CPUs and GPUs, to exploit the data-parallel nature of rendering. The system has an out-of-core data management base layer designed to support a broad variety of rendering algorithms - including volume rendering for visualization purposes. The long-term goal is to facilitate high performance, large-data visualization by taking advantage of the substantial computational capacity of commodity GPUs.

To date, we have completed an initial prototype of this system, submitted a research paper describing initial results, and are beginning to embark on deploying resources to scale the work to a larger collection of GPU-based resources.

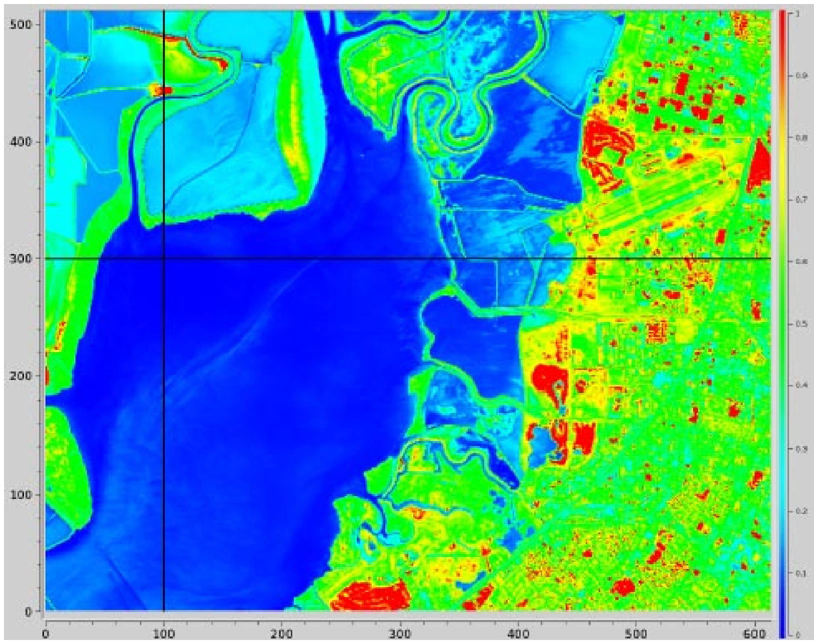

Figure 16: Distance field renderings for feature identification within a two-dimensional function field Here, we show a distance field for a hyperspectral image. The probe is positioned over water (a portion of the San Francisco bay). As expected, other water pixels have low distance values (blue) due to function similarity. Features with functions mapping to high distance (red) include certain types of buildings and land features. The function field in hyperspectral imagery is the integration of radiance versus wavelength. While this example is 2D, the concept is applicable to data of higher dimensionality.

5.5 GPU-Accelerated Query-Driven Visual Data Analysis

Another in VACET's portfolio of basic research projects focuses on leveraging the computational capacity of GPUs to perform index/query operations within the context of high performance query-driven visualization. This effort focuses on two principal areas. The first project aims to develop parallel methods for evaluating queries with specific emphasis to GPU-based approaches. The second project focuses on new methods for extending the power of query-driven analysis methods to adaptive mesh refinement (AMR) data. In this method, query-driven visualization is used to construct multi-temporal visualizations of time-dependent AMR data.

The science stakeholders for this work include John Bell and his team at LBNL who are studying the impact of turbulence on combustion. This work is being performed in collaboration with John Wu (LBNL) of the SciDAC Scientific Data Management Center. Portions of this work are performed in collaboration with John Owens (UCD) of the SciDAC Institute for Ultrascale Visualization.

In this period, we have achieved initial results in both thrust areas and submitted papers presenting early research.

Our preliminary implementation of queries on AMR data creates a derived field indicating refinement level and patch ID; these fields can be included as query conditions as with any other variable. This approach allows us to perform queries at specified levels of AMR refinement, or, alternately, to discover at which level of refinement a query result resides. Our approach directly addresses the challenges associated with visualizing AMR data's dynamic spatial and temporal properties, and is the first-ever work that provides an implementation of query-driven visualization

methods on the GPU, as well as for AMR data. A paper on this research result was submitted to IEEE Visualization 2008. Our GPU-based query-driven visualization technique outperforms existing compressed bitmap indexing techniques during query processing, and offers a tremendous acceleration in performance for query-driven visualization methods. A paper on this research result was submitted to VLDB 2008.

5.6 Comparative Visual Data Analysis

This project impacts a potentially large audience - those who need the ability to perform comparisons on large-scale complex data sets. Nearly all our science stakeholders have requested some sort of capability for comparing data: simulation-simulation, experiment-simulation, and ensemble/parameter studies.

Early work in this space as has focused on creating an initial prototype, implementing it in VisIt, validating with sample stakeholder datasets and submitting a publication outlining early research results.

6 VACET Communication

6.1 VACET Internal Communication Vehicles

Since the VACET team is distributed both geographically and institutionally, we have come to rely heavily on three different forms of internal communication: email, telephone and wiki.

We have established two different email lists: one for the entire VACET team, the other for the VACET EC. These lists are hosted at LBNL and are quite active. Since we are using *Mailman* to provide email list service, we also benefit from having access to browsable archives of email traffic on both lists.

We have established a standing, toll-free conference call line hosted at LBNL for use by team members. The objective is to make it as easy as possible for team members to quickly arrange conference calls. This project-wide service has proven very useful in ad-hoc meetings to address time-critical specific subjects. Examples of such meetings include: organization of a VACET workshop at the 2007 SciDAC program meeting; several meetings aimed at coordinating a VACET-wide response for Science Application Partnership proposals that are part of the Spring 2007 Fusion SciDAC call for proposals.

We have created a project-wide wiki that is hosted at SCI (see `http://www.sci.utah.edu/vacetwiki`. The wiki has both public and authentication-required content areas. Its primary role so far is as a vehicle for posting and sharing documents, collaborative document development, individual stakeholder needs assessment and proposed work tasks, per-project management plans, and so forth. Related, SCI hosts a project-wide *subversion* server that has proven very useful for collaborative, revision-control document and code development.

6.2 VACET Internal Communication Channels

The VACET management has adopted a position of encouraging frequent communication amongst all team members. This project-wide policy has proven to be highly beneficial in terms of achieving: (1) coordination of project-wide activities; (2) helping each team member to feel an active part of a larger, vibrant effort by knowing what everyone else is working on, project objectives, and so forth; (3) fostering an overall sense of unrestricted and open communication amongst all team members.

The VACET Executive Committee holds bi-weekly conference calls. The purpose of these calls is to discuss project management issues, including but not limited to: (1) stakeholder relations; (2) defining and prioritizing VACET technical objectives; (3) inter-institution technical interactions. Minutes of these calls are emailed to the entire team. In addition to the regularly scheduled bi-weekly calls, the EC periodically holds out-of-cycle calls to discuss time critical issues when needed.

The VACET Software Engineering Team holds monthly conference calls. Minutes of the calls are posted on the VACET wiki.

In addition to regular calls, these groups periodically hold ad-hoc meetings on an as-needed basis to address time-critical matters.

In this reporting period, we have held two VACET-wide all-hands meetings. One was held in conjunction with IEEE Visualization 2007 in Sacramento, CA during October 2007. The other was held in April 2008 at the Scientific Computing Institute, University of Utah. The agenda at both meetings is to give project leaders an opportunity to make informal reports of progress, to facilitate face-to-face interaction amongst the geographically distributed team, to introduce new team members to the larger group, to perform inter-project coordination, and to perform global project planning.

6.3 Communication with Other SciDAC Projects

As evidenced by a substantial amount of focused progress in specific science areas in other sections of this document, VACET team members have been very active in reaching out to stakeholder projects. Such outreach includes email, phone calls, and presence at their project meetings. Based upon the early success of such outreach, we expect such frequent and productive communication to continue in the future.

As one of the guiding principles for VACET operations is free and open communication, we have a large amount of "surface area" with many different SciDAC projects: each VACET team member is an "ambassador" for the project. So far, this approach has been hugely successful in fostering broad communication with many different current and potential future stakeholders.

We regularly participate in the project-wide meetings of other SciDAC projects. Examples include:

- Astrophysics: Community Astrophysics Consortium All-Hands Meeting, April 7-9, 2008, Palo Alto, CA. VACET team members gave invited talks describing ongoing joint projects and VACET technology.
- Fusion: FACETS All-Hands Meeting, January 21-22, 2008, Boulder, CO. VACET sent two persons to this meeting; they engaged in several one-on-one meetings with fusion stakeholders to discuss progress to-date and to establish fusion priorities in the area of visual data analysis needs.
- SDM: SDM Center All Hands Meeting, November 28-29, 2007, Seattle, WA. VACET team members attended the workshop, gave presentations, and participated in a mini-workshop focusing on collaboration between SDM and VACET.

6.4 External Communication

Our strategy for outreach and communication to "the world" includes a broad array of activities: web, press, technical publications, and technical presentations.

VACET established a website at `www.vacet.org`. The website, while evolving, is intended to provide a comprehensive view of our Center's activities to a wide audience (science stakeholders, DOE program office, general public).

Press releases provide a valuable way to reach out to a potentially large audience that spans academia, government and industry. One example of such a press release is the November 3, 2006 HPCWire interview with VACET team member Bethel. These types of activities help not only VACET, but also SciDAC and DOE's technology portfolio as a whole.

VACET team members have made presentations in many forums since commencing operations. These are listed in Section 7.5.

VACET team members have a strong record of technical publications. We expect this trend to continue in the future, thereby maintaining VACET's leadership role in the visualization community. These are listed in Section 7.1.

7 Publications, Presentations, Awards and Outreach

7.1 Publications

7.1.1 Peer-Reviewed Journal Articles

1. Scott E. Dillard, Vijay Natarajan, Gunther H. Weber, Valerio Pascucci, and Bernd Hamann. Topology-guided Tessellation of Quadratic Elements. *International Journal of Computational Geometry and Applications (IJCGA) [to appear]*, 2008. LBNL-63771, in press.

2. C. Dietrich, J. Comba, L. Nedel, C. Scheidegger, J. Schreiner, and C. Silva. Improving Mesh Quality of Marching Cubes Using Edge Transformations. *IEEE Transactions on Visualization and Computer Graphics [to appear]*, 2008.

3. Antonio Baptista, Bill Howe, Juliana Freire, David Maier, and Claudio T. Silva. Scientific Exploration in the Era of Ocean Observatories. *Computing in Science & Engineering*, 10(3):53–58, May-June 2008.

4. Erik W. Anderson, James P. Ahrens, Katrin Heitmann, Salman Habib, and Claudio T. Silva. Provenance in Comparative Analysis: A Study in Cosmology. *Computing in Science & Engineering*, 10(3):30–37, May-June 2008.

5. Juliana Freire, David Koop, Emanuele Santos, and Claudio T. Silva. Provenance for Computational Tasks: A Survey. *Computing in Science & Engineering*, 10(3):11–21, May-June 2008.

6. Claudio T. Silva and Joel E. Tohline. Computational Provenance. *Computing in Science & Engineering*, 10(3):9–10, May-June 2008.

7. Joel Daniels, Tilo Ochotta, Linh Ha, and Cláudio T. Silva. Spline-Based Feature Curves from Point-Sampled Geometry. *The Visual Computer*, 2008. To appear.

8. A.R. Sanderson, M.D. Meyer, R.M. Kirby, and C.R. Johnson. A Framework for Exploring Numerical Solutions of Advection-Reaction-Diffusion Equations Using a GPU-Based Approach. *Computing and Visualization in Science*, 2008.

9. E. Wes Bethel, Chris Johnson, Ken Joy, Sean Ahern, Valerio Pascucci, Hank Childs, Jonathan Cohen, Mark Duchaineau, Bernd Hamann, Charles Hansen, Dan Laney, Peter Lindstrom,

Jeremy Meredith, George Ostrouchov, Steven Parker, Claudio Silva, Allen Sanderson, and Xavier Tricoche. SciDAC Visualization and Analytics Center for Enabling Technology. *Journal of Physics Conference Series – SciDAC 2007*, 78, June 2007. LBNL-63542.

10. C. Garth, F. Gerhardt, X. Tricoche, and H. Hagen. Efficient Computation and Visualization of Coherent Structures in Fluid Flow Applications. *IEEE Transactions on Visualization and Computer Graphics*, 13(6):1464–1471, 2007.

11. Gunther H. Weber, Peer-Timo Bremer, and Valerio Pascucci. Topological Landscapes: A Terrain Metaphor for Scientific Data. *IEEE Transactions on Visualization and Computer Graphics (Special Issue: Proceedings of IEEE Visualization 2007)*, 13(6):1416–1423, November/December 2007. LBNL-63763.

12. C. Jones, K.-L. Ma, A. Sanderson, and L. Myers. Visual Interrogation of Gyrokinetic Particle Simulations. *Journal of Physics Conference Series – SciDAC 2007*, 78, June 2007.

13. C. Scheidegger, H. Vo, D. Koop, J. Freire, and C.T. Silva. Querying and creating visualizations by analogy. *IEEE Transactions on Visualization and Computer Graphics*, 13(6):1560–1567, 2007. Winner of the Best Paper Award.

14. I. Wald, H. Friedrich, A. Knoll, and C.D. Hansen. Interactive Isosurface Ray Tracing of Time-Varying Tetrahedral Volumes. *IEEE Transactions on Visualization and Computer Graphics*, 13(6):1727–1734, 2007.

15. Kenneth I. Joy, Mark Miller, Hank Childs, E. Wes Bethel, John Clyne, George Ostrouchov, and Sean Ahern. Frameworks for Visualization at the Extreme Scale. *Journal of Physics Conference Series – SciDAC 2007*, 78, 2007. LBNL-63762.

16. A. Wiebel, X. Tricoche, D. Schneider, Heike Jänicke, and Gerik Scheuermann. Generalized Streak Lines: Analysis and Visualization of Boundary Induced Vortices. *IEEE Transactions on Visualization and Computer Graphics*, 13(6):1735–1742, 2007.

17. Luke Gosink, John C. Anderson, E. Wes Bethel, and Kenneth I. Joy. Variable Interactions in Query-Driven Visualization. *IEEE Transactions on Visualization and Computer Graphics (Proceedings of Visualization 2007)*, 13(6):1400–1407, November/December 2007. LBNL-63524.

18. Attila Gyulassy, Vijay Natarajan, Bernd Hamann, Mark Duchaineau, Valerio Pascucci, Eduardo Bringa, and Andres Higginbotham. Topologically Clean Distance Fields. *IEEE Transactions on Visualization and Computer Graphics (Proceedings of Visualization 2007)*, 13(6):1432–1439, November/December 2007.

19. Sung-Eui Yoon and Peter Lindstrom. Random-Accessible Compressed Triangle Meshes. *IEEE Transactions on Visualization and Computer Graphics (Proceedings of Visualization 2007)*, 13(6):1536–1543, November/December 2007.

20. Valerio Pascucci, Giorgio Scorzelli, Peer-Timo Bremer, and Ajith Mascarenhas. Robust Online Computation of Reeb Graphs: Simplicity and Speed. *ACM Trans. Graph.*, 26(3):58, 2007.

21. P.-T. Bremer, E. Mm Bringa, M. A. Duchaineau, A. G. Ghylassy, D. Laney, A. Mascarenhas, and V. Pascucci. Topological Feature Extraction and Tracking. *Journal of Physics Conference Series – SciDAC 2007*, 78, June 2007.

22. Hank Childs. Architectural Challenges and Solutions for Petascale Postprocessing. *Journal of Physics Conference Series – SciDAC 2007*, 78, June 2007.

23. Guo-Shi Li, Xavier Tricoche, Daniel Weiskopf, and Charles Hansen. Flow Charts: Visualization of Vector Fields on Arbitrary Surfaces. *IEEE Transactions on Visualization and Computer Graphics[accepted]*, 2008.

24. Frank Michel, Eduard Deines, Martin Hering-Bertram, Christoph Garth, and Hans Hagen. Listener-based Analysis of Surface Importance for Acoustic Metrics. *IEEE Transactions on Visualization and Computer Graphics (Special Issue: Proceedings of IEEE Visualization 2007)*, 13(6):1680–1678, November/December 2007.

25. Michael Schlemmer, Manuel Heringer, Florian Morr, Ingrid Hotz, Martin Hering-Bertram, Christoph Garth, Wolfgang Kollmann, Bernd Hamann, and Hans Hagen. Moment invariants for the analysis of 2d flow fields. *IEEE Transactions on Visualization and Computer Graphics (Special Issue: Proceedings of IEEE Visualization 2007)*, 13(6):1743–1750, November/December 2007.

7.1.2 Conference Proceedings

1. Christoph Garth, Xavier Tricoche, Alexander Wiebel, and Ken Joy. On the role of domain-specific knowledge in the visualization of technical flows. In H. Theisel H. Hauser, S. Strassburger, editor, *Proc. Simulation and Visualization 2008 (SimVis '08)*, pages 107–120. SCS Publishing House, February 2008.

2. Christoph Garth, Alexander Wiebel, Xavier Tricoche, Kenneth I. Joy, and Gerik Scheuermann. Lagrangian Visualization of Flow-Embedded Surface Structures. In *Computer Graphics Forum (Proceedings of Eurovis 2008) [to appear]*, 2008.

3. John C. Anderson, Christoph Garth, Mark A. Duchaineau, and Kenneth I. Joy. Discrete Multi-Material Interface Reconstruction for Volume Fraction Data. In *Computer Graphics Forum (Proceedings of Eurovis 2008) [to appear]*, 2008.

4. John Anderson, Luke Gosink, Mark Duchaineau, and Kenneth I. Joy. On Reconstructing Material Interfaces for Visualization. In *Proceedings of the 2008 European Visualization Conference (EuroVis) [to appear]*, 2008.

5. Guo-Shi Li, Xavier Tricoche, and Charles Hansen. Physically-based Dye Advection for Flow Visualization. *Computer Graphics Forum (proceedings of EuroVis 2008) [accepted]*, 2008.

6. Gunther H. Weber, Vincent E. Beckner, Hank Childs, Terry J. Ligocki, Mark Miller, Brian van Straalen, and E. Wes Bethel. Visualization Tools for Adaptive Mesh Refinement Data. In Werner Benger, Rene Heinzel, Wolfgang Kapferer, Wolfram Schoor, Mayank Tyagi, Shalini Venkataraman, and Gunther H. Weber, editors, *Proceedings of the 4th High End Visualization Workshop*, pages 12–25, Berlin, Germany, 2007. Lehmans Media. LBNL-62954.

7. S.G. Parker, S. Boulos, J. Bigler, and A. Robison. RTSL: a Ray Tracing Shading Language. In *In 2007 IEEE/Eurographics Symposium on Interactive Ray Tracing*, 2007. Best Paper Award winner.

8. Gunther H. Weber, Vincent E. Beckner, Hank Childs, Terry J. Ligocki, Mark Miller, Brian van Straalen, and E. Wes Bethel. Visualization of Scalar Adaptive Mesh Refinement Data. In *Numerical Modeling of Space Plasma Flows: Astronum-2007 (Astronomical Society of the Pacific Conference Series)*, volume 385, pages 309–320, 2008.

7.1.3 Invited Articles

1. E.W. Bethel, C.R. Johnson, C. Aragon, Prabhat, O. Rübel, G. Weber, V. Pascucci, H. Childs, P.-T. Bremer, B. Whitlock, S. Ahern, J. Meredith, G. Ostrouchov, K. Joy, B. Hamann, C. Garth, M. Cole, C. Hansen, S. Parker, A. Sanderson, C.T. Silva, and X. Tricoche. DOE's SciDAC Visualization and Analytics Center for Enabling Technologies - Strategy for Petascale Visual Data Analysis Success. *CTWatch Quarterly*, 3(4), 2007.

2. Valerio Pascucci, Peer-Timo Bremer, and Ajith Mascarenhas. Topological Analysis Provides Deeper Insight into Hydrodynamic Instabilities. `http://www.sc.doe.gov/ascr/News/MonthlyNewsRoundup9-07.html`, September 2007.

3. E. Wes Bethel, Chris Johnson, Ken Joy, Sean Ahern, Valerio Pascucci, Hank Childs, Cecilia Aragon, Peer-Timo Bremer, Kathleen Bonnell, Martin Cole, Bernd Hamann, Charles Hansen, Daniel Laney, Ajith Mascarenhas, Jeremy Meredith, George Ostrouchov, Steven Parker, Prabhat, Dave Pugmire, Allen Sanderson, Claudio Silva, Xavier Tricoche, and Gunther Weber. SciDAC's Visualization and Analytics Center for Enabling Technologies. *SciDAC Review (to appear)*, 2008.

7.1.4 Book Chapters

1. Y. Livnat, S. Parker, and C.R. Johnson. Fast Isosurface Extraction Methods for Large Image Data Sets. In Isaac Bankman, editor, *Handbook of Medical Imaging: Processing and Analysis*, chapter 44. Academic Press, 2nd edition edition, 2007. (to appear).

7.1.5 Posters

1. E.W. Bethel, C.R. Johnson, C. Aragon, Prabhat, O. Rübel, G. Weber, V. Pascucci, H. Childs, P.-T. Bremer, A. Mascarenhas, B. Whitlock, S. Ahern, J. Meredith, G. Ostrouchov, K. Joy, B. Hamann, C. Garth, M. Cole, C. Hansen, S. Parker, A. Sanderson, C.T. Silva, and X. Tricoche. DOE SciDAC Visualization and Analytics Center for Enabling Technologies. In *2008 DOE ASCR CS PI Meeting*, Denver, CO, USA, April 2008.

2. Ray Grout, Evatt Hawkes, Jacqueline Chen, Ajith Mascarenhas, Peer-Timo Bremer, and Valerio Pascucci. Analysis of the Relationship Between High Scalar Dissipation Rate Features, Flow, and Combustion. In *WIPP 2008 Combustion Symposium [submitted]*, 2008.

3. Cecilia Aragon, Stephen Bailey, Sarah Poon, Karl Runge, and Rollin Thomas. Sunfall: A Collaborative Visual Analytics System for Astrophysics. In *IEEE Symposium on Visual Analytics Science and Technology (VAST)*, Sacramento, CA, USA, November 2007. Best Poster Award Winner.

7.1.6 Recently Submitted Publications

1. Peer-Timo Bremer, Gunther H. Weber, Marc Day, John Bell, and Valerio Pascucci. Analyzing and Tracking Burning Structures in Lean Premixed Hydrogen Flames. *IEEE Transactions on Visualization and Computer Graphics [submitted]*, 2008.

2. Martin Isenburg, Peter Lindstrom, and Hank Childs. Generating Ghost Data in a Distributed Memory, Streaming Setting. In *Supercomputing 2008 [submitted]*, 2008.

3. John Anderson, Luke Gosink, Mark Duchaineau, and Kenneth I. Joy. An Active Contour Approach to Material Interface Reconstruction. *IEEE Transactions on Visualization and Computer Graphics [submitted]*, 2008.

4. Ajith Mascarenhas, Peer-Timo Bremer, Valerio Pascucci, , Jacqueline Chen, Ray Grout, and Evatt Hawkes. Analysis of High Scalar Dissipation Features in Non-premixed Flames. In *Supercomputing 2008 [submitted]*, 2008.

5. Oliver Rübel, Prabhat, Kesheng Wu, Hank Childs, Jeremy Meredith, Cameron G.R. Geddes, Estelle Cormier-Michel, Sean Ahern, Gunther H. Weber, Peter Messmer, Hans Hagen, Bernd Hamann, and E. Wes Bethel. High Performance Multivariate Visual Data Exploration for Extremely Large Data. In *Supercomputing 2008 [submitted]*, 2008.

6. Luke Gosink, John Wu, Wes Bethel, John Owens, and Kenneth I. Joy. Bin-Hash Indexing: A Parallel GPU-Based Method For Fast Query Processing. In *2008 Conference on Very Large Databases (VLDB) [submitted]*, 2008.

7. Luke Gosink, John Anderson, Wes Bethel, and Kenneth I. Joy. Query-Driven Visualization of Time-Varying Adaptive Mesh Refinement Data. *IEEE Transactions on Visualization and Computer Graphics/Visualization 2008 [submitted]*, 2008.

8. Brian Budge, Tony Bernardin, Jeff Stuart, Shubho Sengupta, John Owens, and Kenneth I. Joy. Out-of-core Data Management for Path Tracing on Hybrid Resources. In *Siggraph, Asia [submitted]*, 2008.

9. Kristin Potter, Joe Kniss, Richard Riesenfeld, and Christopher Johnson. Visualizing Summary Statistics and Uncertainty. *IEEE Transactions on Visualization and Computer Graphics [submitted]*, 2008.

10. Kristin Potter, Jens Krüger, and Christopher Johnson. Towards the Visualization of Multi-dimensional Probabilistic Distribution Data. In *IADIS International Conference on Computer Graphics and Visualization [submitted]*, 2008.

11. Mario Hlawitschka, Christoph Garth, Xavier Tricoche, Gordon Kindlmann, Kenneth I. Joy, and Gerik Scheuermann. Extraction and Visualization of Coherent Fiber Traces in DT-MRI Data. *IEEE Transactions on Visualization and Computer Graphics/Visualization 2008 [submitted]*.

12. Christoph Garth, Hari Krishnan, Xavier Tricoche, Tom Bobach, and Kenneth I. Joy. Generation of Accurate Integral Surfaces Generation of Accurate Integral Surfaces in Time-Dependent Vector Fields. *IEEE Transactions on Visualization and Computer Graphics/Visualization 2008 [submitted]*.

13. Steven P. Callahan and Cláudio T. Silva. Image-Space Acceleration for Direct Volume Rendering of Unstructured Grids using Joint Bilateral Upsampling. *IEEE Transactions on Visualization and Computer Graphics/Visualization 2008 [submitted]*, 2008.

14. Linh K. Ha, P. Thomas Fletcher, Sarang Joshi, and Cláudio T. Silva. Fast Parallel Unbiased Diffeomorphic Atlas Construction on Multi-Graphics Processing Units. *IEEE Transactions on Visualization and Computer Graphics/Visualization 2008 [submitted]*, 2008.

15. C. Dietrich, C. Scheidegger, J. Comba, L. Nedel, and Cláudio T. Silva. Edge Groups: A New Approach to Understanding the Mesh Quality of Marching Methods. *IEEE Transactions on Visualization and Computer Graphics/Visualization 2008 [submitted]*, 2008.

16. Carlos Scheidegger, John Schreiner, Brian Duffy, Hamish Carr, and Cláudio T. Silva. Revisiting Histograms and Isosurface Statistics. *IEEE Transactions on Visualization and Computer Graphics/Visualization 2008 [submitted]*, 2008.

17. David Koop, Carlos E. Scheidegger, Steven P. Callahan, Juliana Freire, and Cláudio T. Silva. VisComplete: Data-driven Suggestions for Visualization Systems. *IEEE Transactions on Visualization and Computer Graphics/Visualization 2008 [submitted]*, 2008.

18. Bill Howe, Peter Lawson, Renee Bellinger, ErikAnderson, Emanuele Santos, Juliana Freire, Carlos Scheidegger, António Baptista, and Cláudio Silva. Data Integration and Visualization for Ocean Observatories. *IEEE Transactions on Visualization and Computer Graphics/Visualization 2008 [submitted]*, 2008.

19. Gilbert A. Preston Erik W. Anderson and Cláudio T. Silva. Visualizing the time-frequency evolution of signal ensembles. *IEEE Transactions on Visualization and Computer Graphics/Visualization 2008 [submitted]*, 2008.

20. Cecilia R. Aragon, Sarah Poon, Greg Aldering, Rollin C. Thomas, and Robert Quimby. Using Visual Analytics to Maintain Situation Awareness in Astrophysics. In *VAST 2008 [submitted]*, 2008.

21. Sarah Poon, Rollin C. Thomas, Cecilia Aragon, and Brian Lee. Virtual Assistants in Collaborative Science Experiments: An Astrophysics Case Study. In *CSCW 2008 [submitted]*, 2008.

22. Hank Childs, Sean Ahern, Jeremy Meredith, John Clyne, and Kenneth I. Joy. Comparative Visualization Using Cross-Mesh Field Evaluations and Derived Quantities. In *Computer Graphics Forum (Proceedings of Eurovis 2008) [submitted]*, 2008.

7.2 Presentations

7.2.1 Invited Presentations

1. E. Wes Bethel. Visual Data Analysis and Data Exploration at the Extreme Scale. In *DOE Office of Science, Advanced Scientific Computing Research Principal Investigator Meeting*, Denver, CO, USA, April 2008.

2. Chris Johnson. Moving Beyond Pretty Pictures. In *IEEE Visualization 2007 (Capstone Speaker)*, Sacramento, CA, USA, October 2007.

3. E. Wes Bethel. Query-Driven Visualization Accelerates Scientific Insight. In *National Science Foundation (NSF) and Cyber-enabled Discovery and Innovation (CDI) Workshop*, Mathematical Sciences Research Institute, Berkeley CA, USA, October 2007.

4. E. Wes Bethel. Occam's Razor and Petascale Visual Data Analysis. In *2007 Falls Creek Falls Conference*, Nashville, TN, USA, October 2007.

5. E. Wes Bethel. Visualization, VACET and the SciDAC Compass Accelerator Project. In *Community Petascale Project for Accelerator Science and Simulation (COMPASS) All-Hands Meeting*, Fermi National Laboratory, Batavia IL, USA, September 2007.

6. Valerio Pascucci. Robust Extraction and Tracking of Topological Features in Scientific Data: State of the Art and Future Challenges. In *Workshop on Feature Extraction and Tracking (Sponsored by the Institute for Ultrascale Visualization)*, UC Davis, Davis CA, USA, August 2007.

7. Hank Childs. Architectural Challenges and Solutions for Petascale Visualization and Analysis. In *2007 Falls Creek Falls Conference*, Nashville, TN, USA, September 2007.

8. Valerio Pascucci. Robust Extraction and Tracking of Topological Features in Scientific Data: State of the Art and Future Challenges. In *Workshop on Feature Extraction and Tracking (Sponsored by the Institute for Ultrascale Visualization)*, UC Davis, Davis CA, USA, August 2007.

9. Hank Childs. VisIt: Visualization and Analysis. In *Computational Astrophysics Consortium SciDAC Center: All Hands Meeting*, Palo Alto, Ca, USA, April 2008.

10. Cecilia Aragon. SpectraVis: Visualization and Analytics. In *Computational Astrophysics Consortium SciDAC Center: All Hands Meeting*, Palo Alto, Ca, USA, April 2008.

11. Cecilia Aragon. Sunfall: Visual Analytics for Astrophysics. In *CAHSI Lecture Series, University of Texas at El Paso*, El Paso, TX, USA, April 2008.

12. Hank Childs. VisIt Overview. In *Scientific Data Management SciDAC Center: All Hands Meeting*, Seattle, Wa, USA, November 2007.

13. David Pugmire, Hank Childs, and Sean Ahern. Parallel Analysis and Visualization on Cray Compute Node Linux. In *DOE Computer Graphics Forum*, Duck, NC, USA, April 2008.

14. Hank Childs. VisIt Update. In *DOE Computer Graphics Forum*, Duck, NC, USA, April 2008.

15. Gunther H. Weber. Visualization Tools for Adaptive Mesh Refinement Data. In *DOE Computer Graphics Forum*, Duck, NC, USA, April 2008.

16. David Pugmire, Hank Childs, and Sean Ahern. Parallel Analysis and Visualization on Cray Compute Node Linux. In *Cray Users Group*, Helsinki, Finland, May 2008.

17. Gunther H. Weber and Peter Nugent. Introduction into VisIt. In *UC Berkeley Astrophysics Computation Discussion Group*, Berkeley, CA, USA, February 2008.

18. Chris Johnson. Putting it All Together: Highlights of Recent National Reports on Computing. In *DOE ASC PI Meeting*, Monterey, CA, USA, February 2008.

19. Chris Johnson. Visual Computing and Imaging: Interdisciplinary Approaches. In *Center for Interdisciplinary Art and Technology*, Salt Lake City, UT, USA, February 2008.

20. Chris Johnson. Large Scale Visual Data Analysis. In *Ultrascale Visualization Workshop*, Reno, NV, USA, November 2007.

21. Luke Gosink, John C. Anderson, E. Wes Bethel, and Kenneth I. Joy. Variable Interactions in Query-Driven Visualization. *IEEE Transactions on Visualization and Computer Graphics (Proceedings of Visualization 2007)*, 13(6):1400–1407, November/December 2007. LBNL-63524.

7.3 Service

7.3.1 Conference Chair

- 2007 IEEE Visualization General Conference – Ken Joy, UC Davis.
- 2007 ASCR Visualization and Analytics Workshop: Chris Johnson, Utah (co-Chair).
- 2007 Ultrascale Visualization Workshop: Chris Johnson, Utah (co-Chair).
- 2008 IEEE Visualization, General Conference: Ken Joy, UC Davis (co-Chair).
- 2008 Volume Graphics Symposium: Valerio Pascucci, LLNL (co-Chair).
- 2008 DOE Computer Graphics Forum, Technical Program Chair: Dave Pugmire, ORNL.

7.3.2 Conference Program Committees

- 2007 IEEE Visualization: Papers Chair, Chuck Hansen, Utah.
- 2007 IEEE Visualization Technical Program: Wes Bethel, LBNL; Valerio Pascucci, LLNL; Claudio Silva, Utah.
- 2007 IEEE Visualization Birds-of-a-Feather Chair: Gunther Weber, LBNL.
- 2007 IEEE Visualization, Local A/V Chair: Gunther Weber, LBNL.
- 2007 SciDAC Program Meeting: Wes Bethel, LBNL.
- 2007 TopoInVis: Ken Joy, UC Davis.
- 2007 Conference on Interactive Techniques in Computer Graphics and Games (I3D 2007): Ken Joy, UC Davis.
- 2007 First Workshop on Knowledge-Assisted Visualization (KAV 2007): Gunther Weber, LBNL.
- 2008 Eurographics: Valerio Pascucci, LLNL.
- 2008 Symposium on Visual Computing: Valerio Pascucci, LLNL.
- 2008 Eurographics Symposium on Parallel Graphics and Visualization: Valerio Pascucci, LLNL.

- 2008 IEEE/ACM Symposium on 3D Data Processing, Visualization and Transmission: Valerio Pascucci, LLNL.
- 2008 IEEE Visualization Papers Committee: Ken Joy, UC Davis; Chuck Hansen, Utah.
- 2008 Volume Graphics Symposium: Chris Johnson, Utah; Claudio Silva, Utah
- 2008 European Visualization Conference (EuroVis): Ken Joy, UC Davis.
- 2008 Dagstuhl Visualization Conference: Ken Joy, UC Davis (co-Organizer).
- 4th IEEE International Conference on e-Science 2008: Chris Johnson, Utah.
- Second ECCOMAS Conference on Computational Vision and Medical Image Processing, 2009: Chris Johnson, Utah.
- International Conference on Computational Science and its Applications 2008: Chris Johnson, Utah.
- VECPAR 2008: Chris Johnson, Utah.

7.3.3 Technical Reviewer

- IEEE Visualization 2007 Technical Program: about 12 from the VACET team served as technical paper reviewers.
- IEEE Transactions on Visualization and Computer Graphics: many of the VACET team serve as technical reviewers for TVCG.
- International Journal of Software and Informatics: Gunther Weber, LBNL.
- Concurrency and Computation: Practice and Experience: Gunther Weber, LBNL.
- Journal of the Earth Simulator: Gunther Weber, LBNL.

7.3.4 Program Review

- Information and Knowledge Sciences Division, Los Alamos National Laboratory: Ken Joy, UC Davis.

7.3.5 External Advisory Board

- NIH Center for Integrative Biomedical Computing, University of Utah: Ken Joy, UC Davis.
- Fundamental and Computational Science Directorate Review Committee, Pacific Northwest National Laboratory, (2007–): Chris Johnson, Utah.
- Scientists and Engineers for America, Advisory Board (2006–): Chris Johnson, Utah.

7.3.6 Editorial Boards

- SIAM Journal of Scientific Computing, Special Issue on Computational Science and Engineering, 2008: Chris Johnson, Utah.
- DOE Office of Advanced Scientific Computing Research Communications Project Editorial Board (2007–): Chris Johnson, Utah.
- SIAM Computational Science and Engineering, Book Series (2003–): Chris Johnson, Utah.
- Computer Graphics Year in Review (2003–): Chris Johnson, Utah.
- Electronic Transactions in Numerical Analysis (2003–): Chris Johnson, Utah.

7.3.7 Workshop Participation

- DOE Modeling and Simulation at the Exacscale for Energy, Ecological Sustainability and Global Security (E3SGS) Town Hall Meetings. About ten different persons from VACET contributed to this effort.

7.4 Awards

1. **IEEE Visualization 2007 – Best Paper Award Winner.** C. Scheidegger, H. Vo, D. Koop, J. Freire, and C.T. Silva. Querying and creating visualizations by analogy. *IEEE Transactions on Visualization and Computer Graphics*, 13(6):1560–1567, 2007. Winner of the Best Paper Award.
2. **IEEE Visualization 2007 – Best Paper Award Nominee.** Luke Gosink, John C. Anderson, E. Wes Bethel, and Kenneth I. Joy. Variable Interactions in Query-Driven Visualization. *IEEE Transactions on Visualization and Computer Graphics (Proceedings of Visualization 2007)*, 13(6):1400–1407, November/December 2007. LBNL-63524.
3. **IEEE Visualization 2007 – Best Paper Award Nominee.** C. Garth, F. Gerhardt, X. Tricoche, and H. Hagen. Efficient Computation and Visualization of Coherent Structures in Fluid Flow Applications. *IEEE Transactions on Visualization and Computer Graphics*, 13(6):1464–1471, 2007.
4. **IEEE Raytracing Symposium 2007 – Best Paper Award Winner.** S.G. Parker, S. Boulos, J. Bigler, and A. Robison. RTSL: a Ray Tracing Shading Language. In *In 2007 IEEE/Eurographics Symposium on Interactive Ray Tracing*, 2007. Best Paper Award winner.
5. **IEEE Symposium on Visual Analytics Science and Technology 2007 – Best Poster Award Winner.** Cecilia Aragon, Stephen Bailey, Sarah Poon, Karl Runge, and Rollin Thomas. Sunfall: A Collaborative Visual Analytics System for Astrophysics. In *IEEE Symposium on Visual Analytics Science and Technology (VAST)*, Sacramento, CA, USA, November 2007. Best Poster Award Winner.
6. **IBM Faculty Award** – Claudio Silva, University of Utah.

7.5 Outreach

- SciDAC 2008 Tutorial [submitted]: VisIt for high performance visual data analysis.
- SciDAC 2008 Tutorial [submitted]: VisTrails for visual data analysis workflow management and provenance.
- SciDAC 2008 Tutorial [submitted]: VACET will contribute to the SDM Center tutorial on FastBit by presenting use cases of employing FastBit in high performance query-driven visual data analysis case studies.
- Congressional outreach. U. S. Congressman Jim Matheson (Utah), who sits on the Science and Energy committee in the U. S. House of Representatives, visited SCI on 4/11/2008. As part of that visit, he dropped in on the VACET All Hands Meeting, held 4/10 and 4/11/2008 at the Scientific Computing Institute at the University of Utah. VACET members explained to Congressman Matheson the important role played by SciDAC in helping to maintain leadership in scientific research, and the crucial role that VACET plays in terms of providing production-quality software infrastructure for gaining scientific insight from large and complex scientific data being produced by simulations and collected from experiments.

Figure 17: Pictured here are VACET team members and U. S. Congressman Jim Matheson at the Scientific Computing Institute, University of Utah, 4/11/2008. Left to right: V. Pascucci (LLNL), S. Ahern (ORNL), W. Bethel (LBNL), Congressman Matheson, C. Johnson (Utah), K. Joy (UC Davis).

8 Resources

8.1 NERSC/LBNL

VACET applied for and received an ERCAP allocation at NERSC for AY08. We were awarded 50,000 CPU hours, approximately 50TB of archival storage, and approximately 1TB on the NERSC global filesystem (GFS). Recently, we requested and received a temporary GFS quota increase to 5TB in order to support "hero-sized" visual data analysis projects in support of our science stakeholders in accelerator and combustion. That effort, conducted jointly with our stakeholders, produced novel research results as well as multiple paper submissions to IEEE Transactions on Visualization and Computer Graphics and Supercomputing 2008.

8.2 LCF/ORNL

At LCF/ORNL, VACET has access to Jaguar through a Director's Discretionary Allocation. Additionally, we have ongoing access to ORNL's visualization cluster.